

美国数学会经典影印系列



Foundations of p -adic Teichmüller Theory

p 进 Teichmüller 理论基础

Shinichi Mochizuki



高等教育出版社

美国数学会经典影印系列



Foundations of Teichmüller Theory

p 进 Teichmüller 理论基础

Shinichi Mochizuki



高等教育出版社·北京

图字：01-2018-2919 号

Foundations of p-adic Teichmüller Theory, by Shinichi Mochizuki,

first published by the American Mathematical Society.

Copyright © 1999 by the American Mathematical Society and International Press. All rights reserved.

This present reprint edition is published by Higher Education Press Limited Company under authority of the American Mathematical Society and is published under license.

Special Edition for People's Republic of China Distribution Only. This edition has been authorized by the American Mathematical Society for sale in People's Republic of China only, and is not for export therefrom.

本书最初由美国数学会于1999年出版，原书名为 *Foundations of p-adic Teichmüller Theory*,

作者为Shinichi Mochizuki。美国数学会和International Press保留原书版权。

原书版权声明：Copyright © 1999 by the American Mathematical Society and International Press。

本影印版由高等教育出版社有限公司经美国数学会独家授权出版。

本版只限于中华人民共和国境内发行。本版经由美国数学会授权仅在中华人民共和国境内销售，不得出口。

p 进 Teichmüller 理论基础

p Jin Teichmüller Lilun Jichu

图书在版编目 (CIP) 数据

p 进 Teichmüller 理论基础 = Foundations of p -adic Teichmüller Theory : 英文 / (日) 望月新一 (Shinichi Mochizuki) 著. —影印本.

—北京：高等教育出版社，2019.1

ISBN 978-7-04-051008-9

I. ① p … II. ①望… III. ①代数曲线—英文
IV. ①O187.1

中国版本图书馆 CIP 数据核字 (2018) 第 275474 号

策划编辑 吴晓丽

责任编辑 吴晓丽

封面设计 张申申

责任印制 赵义民

出版发行 高等教育出版社

社址 北京市西城区德外大街4号

邮政编码 100120

购书热线 010-58581118

咨询电话 400-810-0598

网址 <http://www.hep.edu.cn>

<http://www.hep.com.cn>

网上订购 <http://www.hepmall.com.cn>

<http://www.hepmall.com>

<http://www.hepmall.cn>

印刷 北京中科印刷有限公司

开本 787mm×1092mm 1/16

印张 34.25

字数 850千字

版次 2019年1月第1版

印次 2019年1月第1次印刷

定价 199.00元

本书如有缺页、倒页、脱页等质量问题，

请到所购图书销售部门联系调换

版权所有 侵权必究

[物料号 51008-00]



美国数学会经典影印系列

出版者的话

近年来,我国的科学技术取得了长足进步,特别是在数学等自然科学基础领域不断涌现出一流的研究成果。与此同时,国内的科研队伍与国外的交流合作也越来越密切,越来越多的科研工作者可以熟练地阅读英文文献,并在国际顶级期刊发表英文学术文章,在国外出版社出版英文学术著作。

然而,在国内阅读海外原版英文图书仍不是非常便捷。一方面,这些原版图书主要集中在科技、教育比较发达的大中城市的大型综合图书馆以及科研院所的资料室中,普通读者借阅不甚容易;另一方面,原版书价格昂贵,动辄上百美元,购买也很不方便。这极大地限制了科技工作者对于国外先进科学技术知识的获取,间接阻碍了我国科技的发展。

高等教育出版社本着植根教育、弘扬学术的宗旨服务我国广大科技和教育工作者,同美国数学会(American Mathematical Society)合作,在征求海内外众多专家学者意见的基础上,精选该学会近年出版的数十种专业著作,组织出版了“美国数学会经典影印系列”丛书。美国数学会创建于1888年,是国际上极具影响力的专业学术组织,目前拥有近30000会员和580余个机构成员,出版图书3500多种,冯·诺依曼、莱夫谢茨、陶哲轩等世界级数学大家都是其作者。本影印系列涵盖了代数、几何、分析、方程、拓扑、概率、动力系统所有主要数学分支以及新近发展的数学主题。

我们希望这套书的出版,能够对国内的科研工作者、教育工作者以及青年学生起到重要的学术引领作用,也希望今后能有更多的海外优秀英文著作被介绍到中国。

高等教育出版社

2016年12月

Table of Contents

Table of Contents	vii
Introduction	1
§0. Motivation	
§0.1. The Fuchsian Uniformization	
§0.2. Reformulation in Terms of Metrics	
§0.3. Reformulation in Terms of Indigenous Bundles	
§0.4. Frobenius Invariance and Integrality	
§0.5. The Canonical Real Analytic Trivialization of the Schwarz Torsor	
§0.6. The Frobenius Action on the Schwarz Torsor at the Infinite Prime	
§0.7. Review of the Case of Abelian Varieties	
§0.8. Arithmetic Frobenius Venues	
§0.9. The Classical Ordinary Theory	
§0.10. Intrinsic Hodge Theory	
§1. Overview of the Contents of the Present Book	
§1.1. Major Themes	
§1.2. Atoms, Molecules, and Nilcurves	
§1.3. The \mathcal{MF}^\vee -Object Point of View	
§1.4. The Generalized Notion of a Frobenius Invariant Indigenous Bundle	
§1.5. The Generalized Ordinary Theory	
§1.6. Geometrization	
§1.7. The Canonical Galois Representation	
§1.8. Ordinary Stable Bundles	
§2. Open Problems	
§2.1. Basic Questions	
§2.2. Canonical Curves and Hyperbolic Geometry	
§2.2.1. Review of Kleinian Groups	
§2.2.2. Review of Three-Dimensional Hyperbolic Geometry	
§2.2.3. Rigidity and Density Results	
§2.2.4. QF-Canonical Curves	

§2.2.5. The Case of CM Elliptic Curves	
§2.2.6. The Third Real Dimension as the Frobenius Dimension	
§2.3. Towards an Arithmetic Kodaira-Spencer Theory	
§2.3.1. The Schwarz Torsor as Dual to the Kodaira-Spencer Morphism	
§2.3.2. Arithmetic Resolutions of the Schwarz Torsor	
 Chapter I: Crys-Stable Bundles	87
§0. Introduction	
§1. Definitions and First Properties	
§1.1. Notation Concerning the Underlying Curve	
§1.2. Definition of a Crys-Stable Bundle	
§1.3. Isomorphisms	
§1.4. De Rham Cohomology	
§2. Moduli	
§2.1. Boundedness	
§2.2. Definition of Various Functors	
§2.3. Representability	
§2.4. Radimmersions	
§3. Further Structure	
§3.1. Crystal in Algebraic Spaces	
§3.2. Hodge Morphisms	
§3.3. Clutching Behavior	
§4. Torally Indigenous Bundles	
§4.1. Definitions	
§4.2. Explicit Computation of Monodromy	
§4.3. Moduli and de Rham Cohomology	
§4.4. Clutching Morphisms	
§5. The Universal Torsor of Torally Indigenous Bundles	
§5.1. Notation	
§5.2. Computation	
§5.3. The Case of Dimension One	
 Chapter II: Torally Crys-Stable Bundles in	
 Positive Characteristic	125
§0. Introduction	
§1. The p -Curvature of a Torally Crys-Stable Bundle	
§1.1. Terminology	
§1.2. The p -Curvature at a Marked Point	
§1.3. The Verschiebung Morphism	

- §1.4. Torally Crys-Stable Bundles of Arbitrary Positive Level
- §1.5. The Geometric Connectedness of $\overline{\mathcal{N}}_{g,r}^\rho$
- §1.6. Degenerations of Torally Crys-Stable Bundles of Positive Level
- §2. Nilpotent Connections of Higher Order**
 - §2.1. Higher Order Connections
 - §2.2. De Rham Cohomology Computations
 - §2.3. Versal Families at Infinity
- §3. Mildly Spiked Bundles**
 - §3.1. Definition and First Properties
 - §3.2. De Rham Cohomology Computations
 - §3.3. Deformation Theory

Chapter III: VF-Patterns 171

- §0. Introduction**
- §1. The Moduli Stack Associated to a VF-Pattern**
 - §1.1. Definition of a VF-Pattern
 - §1.2. Construction of Link Stacks
 - §1.3. The Stack Associated to a VF-Pattern
- §2. Affineness Properties**
 - §2.1. A Trivialization of a Certain Line Bundle on $\overline{\mathcal{N}}_{g,r}^\Pi$
 - §2.2. Some Ampleness Results
 - §2.3. Affine Stacks
 - §2.4. Absolute Affineness
 - §2.5. The Connectedness of the Moduli Stack of Curves

Chapter IV: Construction of Examples 199

- §0. Introduction**
- §1. Explicit Computation in the Case $g=1; r=1; p=5$**
 - §1.1. Irreducible Components of Degree Two
 - §1.2. The Case of Radius 1
 - §1.3. Conclusions
- §2. Higher Order Connections and Lubin-Tate Stacks**
 - §2.1. The Projective Line Minus Three Points
 - §2.2. Elliptic Curves
 - §2.3. Lubin-Tate Stacks
- §3. Anabelian Stacks**
 - §3.1. Basic Definitions
 - §3.2. Nondormant Bundles on the Projective Line Minus Three Points
 - §3.3. Explicit Construction of Spiked Data

Pictorial Appendix

Chapter V: Combinatorialization at Infinity of the Stack of Nilcurves 229

§0. Introduction

§1. Statement of Main Results

§2. The Main Theorem

§2.1. The Aphial Case

§2.2. Grafting on Dormant Atoms I: Virtual p-Curvatures

§2.3. Grafting on Dormant Atoms II: Deformation Theory

§2.4. Proof of the Main Theorem

§3. Examples

§3.1. Consequences in the Case $(g, r) = (1, 1)$

§3.2. Explicit Computations

Pictorial Appendix

Chapter VI: The Stack of Quasi-Analytic Self-Isogenies 273

§0. Introduction

§1. Definition of the Stacks $\overline{\mathcal{Q}}_{g,r}^{\Pi}$

§1.1. Epiperfect Schemes

§1.2. The Epiperfect Category

§1.3. Epiperfect Log Schemes

§1.4. The Definition of the Stack of Quasi-Analytic Self-Isogenies

§2. Deformation and Degeneration Properties of the Stacks $\overline{\mathcal{Q}}_{g,r}^{\Pi}$

§2.1. Lifting Properties of $\overline{\mathcal{Q}}_{g,r}^{\Pi}$

§2.2. Representability and Affineness Properties of $\overline{\mathcal{Q}}_{g,r}^{\Pi}$

§2.3. Embeddings of $\overline{\mathcal{Q}}_{g,r}^{\Pi}$

§2.4. The Lattice of Subobjects of S_W

Chapter VII: The Generalized Ordinary Theory 301

§0. Introduction

§1. The Π -Ordinary Locus

§1.1. The Frobenius Action on the Crystalline Cohomology

§1.2. Interpretation of the Condition of Π -Ordinariness

§1.3. Systems of Canonical Modular Frobenius Liftings

§1.4. The Case of Elliptic Curves

§2. The Closure of the Binary Ordinary Locus

§2.1. The Deperfection of the Closure	
§2.2. The Differentials of the Deperfection	
§2.3. The ω -Closedness of the Binary Ordinary Locus	
§3. Existence Results	
§3.1. The Binary Case	
§3.2. The Spiked Case	
§3.3. Frobenius Liftings in the Very Ordinary Case	
Pictorial Appendix	

Chapter VIII: The Geometrization of Binary-Ordinary Frobenius Liftings 345

§0. Introduction	
§1. The General Framework	
§1.1. Canonical Points	
§1.2. The Meaning of “Geometrization”	
§2. The Binary Case	
§2.1. The Associated Differential Formal Group	
§2.2. The Canonical Uniformizing p -divisible Group	
§2.3. Multi-Uniformization by the Group \mathcal{G}_Λ	
§2.4. Canonical Affine Coordinates	
§2.5. Lubin-Tate Geometries	
§2.6. Anabelian Geometries	
§2.7. Deformation of the System of Frobenius Liftings	
§3. Application to Curves and their Moduli	
§3.1. Frobenius Liftings on the Moduli Stack	
§3.2. Frobenius Liftings on the Universal Curve	
Pictorial Appendix	

Chapter IX: The Geometrization of Spiked Frobenius Liftings . . . 397

§0. Introduction	
§1. The Formal Uniformizing \mathcal{MF}^∇-Object	
§1.1. The Objects in Question	
§1.2. The Strong Portion of the Uniformization	
§1.3. The Strong Portion of the Mantle	
§1.4. The Renormalized Frobenius Pull-back of the Mantle	
§1.5. Hodge Subspaces	
§2. Associated Galois Representations	
§2.1. The Strictly Weak Pair of Frobenius Liftings over the Strong Perfection	
§2.2. The Associated Non-affine Geometry	
§2.3. Construction of the Galois Mantle: The Spiked Case	

- §2.4. Discussion of the Resulting Spiked Geometry
- §2.5. Construction of the Galois Mantle:
The Binary-Ordinary Case
- §3. Application to Curves and their Moduli
 - §3.1. Frobenius Liftings on the Moduli Stack
 - §3.2. Frobenius Liftings on the Universal Curve

Pictorial Appendix

Chapter X: Representations of the Fundamental Group of the Curve	463
§0. Introduction	
§1. The Binary-Ordinary Case	
§1.1. The Formal \mathcal{MF}^∇ -Object	
§1.2. The Crystalline Induced Representation	
§1.3. The Lubin-Tate Case	
§1.4. Relation to the Profinite Teichmüller Group	
§2. The Very Ordinary Spiked Case	
§2.1. The Formal \mathcal{MF}^∇ -Object	
§2.2. The Crystalline Induced Representation	
§2.3. Relation to the Profinite Teichmüller Group	
§3. Conclusion	
 Appendix: Ordinary Stable Bundles on a Curve	 493
§0. Introduction	
§1. The Algebraic Theory	
§1.1. Basic Definitions	
§1.2. Moduli	
§2. The Complex Theory	
§2.1. Unitary Representations of the Fundamental Group	
§2.2. The Kähler Approach	
§3. The Ordinary p-adic Theory	
§3.1. Crystals of Bundles with Connection	
§3.2. Frobenius Actions	
§3.3. The Ordinary Case	
§3.4. Canonical Coordinates via the Weil Conjectures	
 Bibliography	 519
Index	525

Introduction

§0. Motivation

The goal of the present work is to lay the foundations for a theory of *uniformization of p -adic hyperbolic curves and their moduli*. On the one hand, this theory generalizes the Fuchsian and Bers uniformizations of complex hyperbolic curves and their moduli to nonarchimedean places. It is for this reason that we shall often refer to the theory of this book as *p -adic Teichmüller theory*, for short. On the other hand, the theory of this book may be regarded as a fairly precise hyperbolic analogue of the Serre-Tate theory of ordinary abelian varieties and their moduli. Since this theory of uniformization of p -adic hyperbolic curves and their moduli was already initiated in [Mzk1], and, in some sense, the present work is a continuation and generalization of [Mzk1], we would like to take the opportunity in this § to try to bridge the gap for the reader between the classical uniformization of a hyperbolic Riemann surface that one studies in an undergraduate complex analysis course and the point of view espoused in the present work.

§0.1. The Fuchsian Uniformization

Let X be a *hyperbolic* algebraic curve over \mathbb{C} , the field of complex numbers. By this, we mean that X is obtained by removing r points from a smooth, proper, connected algebraic curve of genus g (over \mathbb{C}), where $2g - 2 + r > 0$. We shall refer to (g, r) as the *type* of X . Then it is well-known that to X , one can associate in a natural way a Riemann surface \mathbf{X} whose underlying point set is $X(\mathbb{C})$. We shall refer to Riemann surfaces \mathbf{X} obtained in this way as “hyperbolic of finite type.”

Now perhaps the most fundamental *arithmetic* – read “arithmetic at the infinite prime” – fact known about the *algebraic* curve X is that \mathbf{X} admits a uniformization by the upper half plane \mathbf{H} :

$$\mathbf{H} \rightarrow \mathbf{X}$$

For convenience, we shall refer to this uniformization of \mathbf{X} in the following as the *Fuchsian uniformization* of \mathbf{X} . Put another way, the uniformization theorem quoted above asserts that the universal covering space $\tilde{\mathbf{X}}$ of \mathbf{X} (which itself has the natural structure of a Riemann surface) is holomorphically isomorphic to the upper half plane $\mathbf{H} = \{z \in \mathbf{C} \mid \text{Im}(z) > 0\}$. This fact was “familiar” to many mathematicians as early as the mid-nineteenth century, but was only proven rigorously much later by K  be.

The fundamental thrust of [Mzk1] and the present work is to generalize the Fuchsian uniformization to the p -adic context.




Fig. 1: The Fuchsian Uniformization

At this point, the reader might be moved to interject: But hasn’t this already been achieved decades ago by Mumford in [Mumf2]? In fact, however, Mumford’s construction gives rise to a p -adic analogue *not of the Fuchsian uniformization, but rather of the Schottky uniformization* of a complex hyperbolic curve. Even in the complex case, the Schottky uniformization is an entirely different sort of uniformization – both geometrically and arithmetically – from the Fuchsian uniformization: for instance, its periods are holomorphic, whereas the periods that occur for the Fuchsian uniformization are only real analytic. This phenomenon manifests itself in the nonarchimedean context in the fact that the construction of [Mumf2] really has nothing to do with a fixed prime number “ p ,” and in fact, takes place entirely in the formal analytic category. In particular, the theory of [Mumf2] has nothing to do with “Frobenius.” By contrast, the theory of the present book depends very much on the choice of a prime “ p ,” and makes essential use of the “action of Frobenius.” Another difference between the theory of [Mumf2] and the theory of the present book is that [Mumf2] only addresses the case of curves whose “reduction modulo p ” is totally degenerate, whereas the theory of the present book applies to curves

whose reduction modulo p is only assumed to be “sufficiently generic.” Thus, at any rate, the theory of the present book is entirely different from and has little directly to do with the theory of [Mumf2].

§0.2. Reformulation in Terms of Metrics

Unfortunately, if one sets about trying to generalize the Fuchsian uniformization $\mathbf{H} \rightarrow \mathbf{X}$ to the p -adic case in any sort of naive, literal sense, one immediately sees that one runs into a multitude of apparently insurmountable difficulties. Thus, it is natural to attempt to recast the Fuchsian uniformization in a more universal form, a form more amenable to relocation from the archimedean to the nonarchimedean world.

One natural candidate that arises in this context is the notion of a metric – more precisely, the notion of a *real analytic Kähler metric*. For instance, the upper half plane admits a natural such metric, namely, the metric given by

$$\frac{dx^2 + dy^2}{y^2}$$

(where $z = x + iy$ is the standard coordinate on \mathbf{H}). Since this metric is invariant with respect to all holomorphic automorphisms of \mathbf{H} , it induces a natural metric on $\tilde{\mathbf{X}} \cong \mathbf{H}$ which is independent of the choice of isomorphism $\tilde{\mathbf{X}} \cong \mathbf{H}$ and which descends to a metric $\mu_{\mathbf{X}}$ on \mathbf{X} .

Having constructed the canonical metric $\mu_{\mathbf{X}}$ on \mathbf{X} , we first make the following observation:

There is a general theory of canonical coordinates associated to a real analytic Kähler metric on a complex manifold.

(See, e.g., [Mzk1], Introduction, §2, for more technical details.) Moreover, the canonical coordinate associated to the metric $\mu_{\mathbf{X}}$ is precisely the coordinate obtained by pulling back the standard coordinate “ z ” on the unit disc via any holomorphic isomorphism of $\tilde{\mathbf{X}} \cong \mathbf{H}$ with the unit disc. Thus, in other words, passing from $\mathbf{H} \rightarrow \tilde{\mathbf{X}}$ to $\mu_{\mathbf{X}}$ is a “faithful operation,” i.e., one doesn’t really lose any information.

Next, let us make the following observation: Let $\mathcal{M}_{g,r}$ denote the moduli stack of smooth r -pointed algebraic curves of genus g over \mathbf{C} .

If we order the points that were removed from the compactification of X to form X , then we see that X defines a point $[X] \in \mathcal{M}_{g,r}(\mathbf{C})$. Moreover, it is elementary and well-known that the cotangent space to $\mathcal{M}_{g,r}$ at $[X]$ can be written in terms of square differentials on X . Indeed, if, for simplicity, we restrict ourselves to the case $r = 0$, then this cotangent space is naturally isomorphic to $Q \stackrel{\text{def}}{=} H^0(X, \omega_{X/\mathbf{C}}^{\otimes 2})$ (where $\omega_{X/\mathbf{C}}$ is the algebraic coherent sheaf of differentials on X). Then the observation we would like to make is the following: Reformulating the Fuchsian uniformization in terms of the metric μ_X allows us to “push-forward” μ_X to obtain a canonical real analytic Kähler metric μ_M on the complex analytic stack $\mathbf{M}_{g,r}$ associated to $\mathcal{M}_{g,r}$ by the following formula: if $\theta, \psi \in Q$, then

$$\langle \theta, \psi \rangle \stackrel{\text{def}}{=} \int_X \frac{\theta \cdot \bar{\psi}}{\mu_X}$$

(Here, $\bar{\psi}$ is the complex conjugate differential to ψ , and the integral is well-defined because the integrand is the quotient of a $(2,2)$ -form by a $(1,1)$ -form, i.e., the integrand is itself a $(1,1)$ -form.)

This metric on $\mathbf{M}_{g,r}$ is called the *Weil-Petersson metric*. It is known that

The canonical coordinates associated to the Weil-Petersson metric coincide with the so-called Bers coordinates on $\widetilde{\mathbf{M}}_{g,r}$ (the universal covering space of $\mathbf{M}_{g,r}$).

The Bers coordinates define an anti-holomorphic embedding of $\widetilde{\mathbf{M}}_{g,r}$ into the complex affine space associated to Q . We refer to the Introduction of [Mzk1] for more details on this circle of ideas.

At any rate, in summary, we see that much that is useful can be obtained from this reformulation in terms of metrics. However, although we shall see later that the reformulation in terms of metrics is not entirely irrelevant to the theory that one ultimately obtains in the p -adic case, nevertheless this reformulation is still not sufficient to allow one to effect the desired translation of the Fuchsian uniformization into an analogous p -adic theory.

§0.3. Reformulation in Terms of Indigenous Bundles

It turns out that the “missing link” necessary to translate the Fuchsian uniformization into an analogous p -adic theory was provided by Gunning ([Gunning]) in the form of the notion of an *indigenous bundle*. The basic idea is as follows: First recall that the group $\text{Aut}(\mathbf{H})$ of holomorphic automorphisms of the upper half plane may be identified (by thinking about linear fractional transformations) with $\text{PSL}_2(\mathbf{R})^0$ (where the superscripted “0” denotes the connected component of the identity). Moreover, $\text{PSL}_2(\mathbf{R})^0$ is naturally contained inside $\text{PGL}_2(\mathbf{C}) = \text{Aut}(\mathbf{P}_{\mathbf{C}}^1)$. Let $\Pi_{\mathbf{X}}$ denote the (topological) fundamental group of \mathbf{X} (where we ignore the issue of choosing a base-point since this will be irrelevant for what we do). Then since $\Pi_{\mathbf{X}}$ acts naturally on $\tilde{\mathbf{X}} \cong \mathbf{H}$, we get a natural representation

$$\rho_{\mathbf{X}} : \Pi_{\mathbf{X}} \rightarrow \text{PGL}_2(\mathbf{C}) = \text{Aut}(\mathbf{P}_{\mathbf{C}}^1)$$

which is well-defined up to conjugation by an element of $\text{Aut}(\mathbf{H}) \subseteq \text{Aut}(\mathbf{P}_{\mathbf{C}}^1)$. We shall henceforth refer to $\rho_{\mathbf{X}}$ as the *canonical representation associated to \mathbf{X}* . Thus, $\rho_{\mathbf{X}}$ gives us an action of $\Pi_{\mathbf{X}}$ on $\mathbf{P}_{\mathbf{C}}^1$, hence a diagonal action on $\tilde{\mathbf{X}} \times \mathbf{P}_{\mathbf{C}}^1$. If we form the quotient of this action of $\Pi_{\mathbf{X}}$ on $\tilde{\mathbf{X}} \times \mathbf{P}_{\mathbf{C}}^1$, we obtain a \mathbf{P}^1 -bundle over $\tilde{\mathbf{X}}/\Pi_{\mathbf{X}} = \mathbf{X}$ which automatically algebraizes to an algebraic \mathbf{P}^1 -bundle $P \rightarrow X$ over X . (For simplicity, think of the case $r = 0$!)

In fact, $P \rightarrow X$ comes equipped with more structure. First of all, note that the trivial \mathbf{P}^1 -bundle $\tilde{\mathbf{X}} \times \mathbf{P}_{\mathbf{C}}^1 \rightarrow \tilde{\mathbf{X}}$ is equipped with the trivial connection. (Note: here we use the “Grothendieck definition” of the notion of a connection on a \mathbf{P}^1 -bundle: i.e., an isomorphism of the two pull-backs of the \mathbf{P}^1 -bundle to the first infinitesimal neighborhood of the diagonal in $\tilde{\mathbf{X}} \times \tilde{\mathbf{X}}$ which restricts to the identity on the diagonal $\tilde{\mathbf{X}} \subseteq \tilde{\mathbf{X}} \times \tilde{\mathbf{X}}$.) Moreover, this trivial connection is clearly fixed by the action of $\Pi_{\mathbf{X}}$, hence descends and algebraizes to a connection ∇_P on $P \rightarrow X$. Finally, let us observe that we also have a section $\sigma : X \rightarrow P$ given by descending and algebraizing the section $\tilde{\mathbf{X}} \rightarrow \tilde{\mathbf{X}} \times \mathbf{P}_{\mathbf{C}}^1$ whose projection to the second factor is given by $\tilde{\mathbf{X}} \cong \mathbf{H} \subseteq \mathbf{P}_{\mathbf{C}}^1$. This section is referred to as the *Hodge section*. If we differentiate σ by means of ∇_P , we obtain a *Kodaira-Spencer morphism* $\tau_{X/\mathbf{C}} \rightarrow \sigma^* \tau_{P/X}$ (where “ $\tau_{A/B}$ ” denotes the relative tangent bundle of A over B). It is easy to see that this Kodaira-Spencer morphism is necessarily an isomorphism.




Fig. 2: The Construction of the Canonical Indigenous Bundle

This triple of data $(P \rightarrow X, \nabla_P, \sigma)$ is the prototype of what Gunning refers to as an *indigenous bundle*. We shall refer to this specific $(P \rightarrow X, \nabla_P)$ (one doesn't need to specify σ since σ is uniquely determined by the property that its Kodaira-Spencer morphism is an isomorphism) as the *canonical indigenous bundle*. More generally, an *indigenous bundle* on X (at least in the case $r = 0$) is any \mathbf{P}^1 -bundle $P \rightarrow X$ with connection ∇_P such that $P \rightarrow X$ admits a section (necessarily unique) whose Kodaira-Spencer morphism is an isomorphism. (In the case $r > 0$, it is natural to introduce log structures in order to make a precise definition.)

Note that the notion of an indigenous bundle has the virtue of being *entirely algebraic* in the sense that at least as an object, the canonical indigenous bundle $(P \rightarrow X, \nabla_P)$ exists in the algebraic category. In fact, the space of indigenous bundles forms a torsor over the vector space Q of quadratic differentials on X (at least for $r = 0$). Thus,

The issue of which point in this affine space of indigenous bundles on X corresponds to the canonical indigenous bundle is a deep arithmetic issue, but the affine space itself can be defined entirely algebraically.

One aspect of the fact that the notion of an indigenous bundle is entirely algebraic is that indigenous bundles can, in fact, be defined over $\mathbf{Z}[\frac{1}{2}]$, and in particular, over \mathbf{Z}_p (for p odd). In [Mzk1], Chapter I, a fairly complete theory of indigenous bundles in the p -adic case (analogous to the complex theory of [Gunning]) is worked out. To summarize, indigenous bundles are closely related to projective structures and Schwarzian derivatives on X . Moreover, the underlying \mathbf{P}^1 -bundle $P \rightarrow X$ is always the same (for all indigenous bundles on X), i.e., the choice of connection ∇_P determines the isomorphism class of the indigenous bundle. We refer the reader to [Mzk1], Chapter I, for more details. (Note: Although the detailed theory of [Mzk1], Chapter I, is philosophically very relevant to the theory of the present book, most of this theory is technically and logically unnecessary for reading the present book.)

At any rate, to summarize, the introduction of indigenous bundles allows one to consider the Fuchsian uniformization as being embodied by an object – the canonical indigenous bundle – which exists in the algebraic category, but which, compared to other indigenous bundles, is somehow “special.” In the following, we would like to analyze the sense in which the canonical indigenous bundle is special, and to show how this sense can be translated immediately into the p -adic context. Thus, we see that

The search for a p -adic theory analogous to the theory of the Fuchsian uniformization can be reinterpreted as the search for a notion of “canonical p -adic indigenous bundle” which is special in a sense precisely analogous to the sense in which the canonical indigenous bundle arising from the Fuchsian uniformization is special.

§0.4. Frobenius Invariance and Integrality

In this subsection, we explore in greater detail the issue of what precisely makes the canonical indigenous bundle (in the complex case) so special, and note in particular that a properly phrased characterization of the canonical indigenous bundle (in the complex case) translates very naturally into the p -adic case.

First, let us observe that in global discussions of motives over a number field, it is natural to think of the operation of complex conjugation as a sort of “Frobenius at the infinite prime” (cf. [BK], §5). In fact, in such discussions, complex conjugation is often denoted by “ Fr_∞ .” Next, let us observe that one special property of the canonical indigenous bundle is that its monodromy representation (i.e., the “canonical representation” $\rho_X : \Pi_X \rightarrow \mathrm{PGL}_2(\mathbf{C})$) is *real-valued*, i.e., takes its values in $\mathrm{PGL}_2(\mathbf{R})$. Another way to put this is to say that the canonical indigenous bundle is Fr_∞ -invariant, i.e.,

The canonical indigenous bundle on a hyperbolic curve is invariant with respect to the Frobenius at the infinite prime.

Unfortunately, as is observed in [Falt2], this property of having real monodromy is not sufficient to characterize the canonical indigenous bundle completely. That is to say, the indigenous bundles with real monodromy form a discrete subset of the space of indigenous bundles on the given curve X , but this discrete subset consists (in general) of more than one element.

Let us introduce some notation. Let $\mathcal{M}_{g,r}$ be the stack of r -pointed smooth curves of genus g over \mathbf{C} . Let $\mathcal{S}_{g,r}$ be the stack of such curves equipped with an indigenous bundle. Then there is a natural projection morphism $\mathcal{S}_{g,r} \rightarrow \mathcal{M}_{g,r}$ (given by forgetting the indigenous bundle) which exhibits $\mathcal{S}_{g,r}$ as an affine torsor on $\mathcal{M}_{g,r}$ over the vector bundle $\Omega_{\mathcal{M}_{g,r}/\mathbf{C}}$ of differentials on $\mathcal{M}_{g,r}$. We shall refer to this torsor $\mathcal{S}_{g,r} \rightarrow \mathcal{M}_{g,r}$ as the *Schwarz torsor*.

Let us write \mathcal{S}_X for the restriction of the Schwarz torsor $\mathcal{S}_{g,r} \rightarrow \mathcal{M}_{g,r}$ to the point $[X] \in \mathcal{M}_{g,r}(\mathbf{C})$ defined by X . Thus, \mathcal{S}_X is an affine complex space of dimension $3g - 3 + r$. Let $\mathcal{R}_X \subseteq \mathcal{S}_X$ be the set of indigenous bundles with real monodromy. As observed in [Falt2], \mathcal{R}_X is a discrete subset of \mathcal{S}_X . Now let $\mathcal{S}'_X \subseteq \mathcal{S}_X$ be the subset of indigenous bundles $(P \rightarrow X, \nabla_P)$ with the following property:

(*) The associated monodromy representation $\rho : \Pi_X \rightarrow \mathrm{PGL}_2(\mathbf{C})$ is injective and its image Γ is a *quasi-Fuchsian* group. Moreover, if $\Omega \subseteq \mathbf{P}^1(\mathbf{C})$ is the domain of discontinuity of Γ , then Ω/Γ is a disjoint union of two Riemann surfaces of type (g, r) .

(Roughly speaking, a “quasi-Fuchsian group” is a discrete subgroup of $\mathrm{PGL}_2(\mathbf{C})$ whose domain of discontinuity Ω (i.e., the set of points of $\mathbf{P}^1(\mathbf{C})$ at which Γ acts discontinuously) is a disjoint union of two topological open discs, separated by a topological circle. We refer to [Mask] for more details on the theory of quasi-Fuchsian groups.)

It is known that \mathcal{S}'_X is a bounded ([Gard], p. 99, Lemma 6), open (cf. the discussion of §5 of [Thur]) subset of \mathcal{S}_X (in the complex analytic topology). Moreover, since a quasi-Fuchsian group with real monodromy acts discretely on the upper half plane (see, e.g., [Shi], Chapter I, Proposition 1.8), it follows immediately that such a quasi-Fuchsian group is Fuchsian. Put another way, we have that:

The intersection $\mathcal{R}_X \cap \mathcal{S}'_X \subseteq \mathcal{S}_X$ is the set consisting of the single point corresponding to the canonical indigenous bundle.

It is this characterization of the canonical indigenous bundle that we will seek to translate into the p -adic case.

To translate the above characterization, let us first recall the point of view of Arakelov theory which states, in effect, that \mathbf{Z}_p -integral structures (on say, an affine space over \mathbf{Q}_p) correspond to closures of bounded open subsets (of, say, an affine space over \mathbf{C}). Thus, from this point of view, one may think of S'_X as defining a natural integral structure (in the sense of Arakelov theory) on the complex affine space S_X . Thus, from this point of view, one arrives at the following characterization of the canonical indigenous bundle:

The canonical indigenous bundle is the unique indigenous bundle which is integral (in the Arakelov sense) and Frobenius invariant (i.e., has monodromy which is invariant with respect to complex conjugation).

This gives us at last an answer to the question posed earlier: How can one characterize the canonical indigenous bundle in the complex case in such a way that the characterization carries over word for word to the p -adic context? In particular, it gives rise to the following conclusion:

The proper p -adic analogue of the theory of the Fuchsian and Bers uniformizations should be a theory of \mathbf{Z}_p -integral indigenous bundles that are invariant with respect to some natural action of the Frobenius at the prime p .

This conclusion constitutes the fundamental philosophical basis underlying the theory of this book. In [Mzk1], this philosophy was partially realized in the sense that *certain* \mathbf{Z}_p -integral Frobenius indigenous bundles were constructed. The theory of [Mzk1] will be reviewed later (in §0.9). The goal of the present book, by contrast, is to lay the foundations for a general theory of *all* \mathbf{Z}_p -integral Frobenius indigenous bundles and to say as much as is possible in as much generality as is possible concerning such bundles.

§0.5. The Canonical Real Analytic Trivialization of the Schwarz Torsor

In this subsection, we would like to take a closer look at the *Schwarz torsor* $S_{g,r} \rightarrow \mathcal{M}_{g,r}$. For general g and r , this affine torsor $S_{g,r} \rightarrow \mathcal{M}_{g,r}$ does not admit any algebraic or holomorphic sections. Indeed, this affine torsor defines a class in $H^1(\mathcal{M}_{g,r}, \Omega_{\mathcal{M}_{g,r}/\mathbf{C}})$ which is the Hodge-theoretic first Chern class of a certain ample line bundle \mathcal{L} on $\mathcal{M}_{g,r}$. (See [Mzk1], Chapter I, §3, especially Theorem 3.4, for more details on this Hodge-theoretic Chern class and Chapter III, Proposition 2.2, of the present work for a proof of ampleness.) Put another way, $S_{g,r} \rightarrow \mathcal{M}_{g,r}$ is the torsor of (algebraic) connections on the line bundle \mathcal{L} . However, the map that assigns to X the canonical indigenous bundle on X defines a *real analytic* section

$$s_{\mathbf{H}} : \mathcal{M}_{g,r}(\mathbf{C}) \rightarrow \mathcal{S}_{g,r}(\mathbf{C})$$

of this torsor.

The first and most important goal of the present subsection is to remark that

The single object $s_{\mathbf{H}}$ essentially embodies the entire uniformization theory of complex hyperbolic curves and their moduli.

Indeed, $s_{\mathbf{H}}$ by its very definition contains the data of “which indigenous bundle is canonical,” hence already may be said to embody the Fuchsian uniformization. Next, we observe that $\bar{\partial}s_{\mathbf{H}}$ is equal to the Weil-Petersson metric on $\mathcal{M}_{g,r}$ (see [Mzk1], Introduction, Theorem 2.3 for more details). Moreover, (as is remarked in Example 2 following Definition 2.1 in [Mzk1], Introduction, §2) since the canonical coordinates associated to a real analytic Kähler metric are obtained by essentially integrating (in the “sense of anti- $\bar{\partial}$ -ing”) the metric, it follows that (a certain appropriate restriction of) $s_{\mathbf{H}}$ “is” essentially the Bers uniformization of Teichmüller space. Thus, as advertised above, the single object $s_{\mathbf{H}}$ stands at the very center of the uniformization theory of complex hyperbolic curves and their moduli.

In particular, it follows that we can once again reinterpret the fundamental issue of trying to find a p -adic analogue of the Fuchsian uniformization as the issue of *trying to find a p -adic analogue of the section $s_{\mathbf{H}}$* . That is to say, the torsor $\mathcal{S}_{g,r} \rightarrow \mathcal{M}_{g,r}$ is, in fact, defined over $\mathbf{Z}[\frac{1}{2}]$, hence over \mathbf{Z}_p (for p odd). Thus, forgetting for the moment that it is not clear precisely what p -adic category of functions corresponds to the real analytic category at the infinite prime, one sees that

One way to regard the search for a p -adic Fuchsian uniformization is to regard it as the search for some sort of canonical p -adic analytic section of the torsor $\mathcal{S}_{g,r} \rightarrow \mathcal{M}_{g,r}$.

In this context, it is thus natural to refer to $s_{\mathbf{H}}$ as the *canonical arithmetic trivialization of the torsor $\mathcal{S}_{g,r} \rightarrow \mathcal{M}_{g,r}$ at the infinite prime*.

Finally, let us observe that this situation of a torsor corresponding to the Hodge-theoretic first Chern class of an ample line bundle, equipped with a canonical real analytic section occurs not only over $\mathcal{M}_{g,r}$, but over any individual hyperbolic curve X (say, over \mathbf{C}), as well. Indeed, let $(P \rightarrow X, \nabla_P)$ be the canonical indigenous bundle on X . Let $\sigma : X \rightarrow P$ be its Hodge section. Then by [Mzk1], Chapter I, Proposition 2.5, it follows that the $T \stackrel{\text{def}}{=} P - \sigma(X)$ has the structure of an $\omega_{X/\mathbf{C}}$ -torsor over X . In fact, one can say more: namely, this torsor is the Hodge-theoretic first Chern class corresponding to the ample

line bundle $\omega_{X/\mathbb{C}}$. Moreover, if we compose the morphism $\tilde{X} \cong \mathbb{H} \subseteq \mathbb{P}_{\mathbb{C}}^1$ used to define σ with the standard complex conjugation morphism on $\mathbb{P}_{\mathbb{C}}^1$, we obtain a new Π_X -equivariant $\tilde{X} \rightarrow \mathbb{P}_{\mathbb{C}}^1$ which descends to a real analytic section $s_X : X(\mathbb{C}) \rightarrow T(\mathbb{C})$. Just as in the case of $\mathcal{M}_{g,r}$, it is easy to compute (cf. the argument of [Mzk1], Introduction, Theorem 2.3) that $\bar{\partial}s_X$ is equal to the canonical hyperbolic metric μ_X . Thus, just as in the case of $\mathcal{M}_{g,r}$, s_X essentially “is” the Fuchsian uniformization of X .

§0.6. The Frobenius Action on the Schwarz Torsor at the Infinite Prime

In the last few subsections, we have guided the reader through the following metamorphosis of ideas:

Fuchsian Uniformization \longrightarrow Canonical Metrics
 \longrightarrow Frobenius Invariant Integral Indigenous Bundles
 \longrightarrow Canonical Real Analytic Section of the Schwarz Torsor

Each step in this metamorphosis was intended to aid in weeding out those inessential aspects of the uniformization theory of hyperbolic curves and their moduli that are peculiar to the complex case and thereby to render this theory in a form that translates naturally into the p -adic context. In this final subsection on complex Teichmüller theory, we would like to construct a single object – namely, a Frobenius action on the (integral portion of) the Schwarz torsor at the infinite prime – that we feel ties together all the ideas (metrics, Frobenius invariance and integrality, canonical section) discussed so far. Moreover,

This Frobenius action serves to further solidify the discussion of §0.4 in that it gives us an explicit action (of Frobenius) with respect to which the canonical indigenous bundle is invariant.

This will make the transition to the p -adic case more natural since in the p -adic case (unlike the complex case) there is no direct construction of the canonical indigenous bundle “on an individual curve:” that is, the only way to construct the canonical indigenous bundle on an individual curve is to first construct the canonical p -adic trivialization of the Schwarz torsor over the moduli space, and then restrict this trivialization to the point corresponding to the individual curve in question.

We begin with the complex analytic Schwarz torsor $\mathbf{S}_{\mathbf{g},\mathbf{r}} \rightarrow \mathbf{M}_{\mathbf{g},\mathbf{r}}$. (In this following, we shall use boldface letters to denote the complex

analytic versions of algebraic objects.) Recall that in §0.4, we defined for each $[\mathbf{X}] \in \mathbf{M}_{\mathbf{g},r}$, an open bounded subset $\mathbf{S}'_{\mathbf{X}} \subseteq \mathbf{S}_{\mathbf{X}}$. It follows easily from complex Teichmüller theory (cf. the discussion of §5 of [Thur]) that these open subsets “glue together” to form one big open subset

$$\mathbf{S}'_{\mathbf{g},r} \subseteq \mathbf{S}_{\mathbf{g},r}$$

which, in the spirit of the discussion of §0.4, we may think of as defining an integral structure on $\mathbf{S}_{\mathbf{g},r}$.

As follows from the discussion of §5 of [Thur], one may think of $\mathbf{S}'_{\mathbf{g},r}$ as the space of all isomorphism classes of the following data:

- (1) a quasi-Fuchsian group $\Gamma \in \mathrm{PGL}_2(\mathbf{C})$ with the property that if $\Omega \subseteq \mathbf{P}^1(\mathbf{C})$ is the domain of discontinuity of Γ , then Ω/Γ is a disjoint union of two Riemann surfaces of type (g, r) ;
- (2) a choice of one of the connected components $R \subseteq \Omega/\Gamma$; we shall refer to this component as the *distinguished component*.

The utility of this point of view is that it allows one to see things as follows: The projection map $\pi_1 : \mathbf{S}'_{\mathbf{g},r} \subseteq \mathbf{S}_{\mathbf{g},r} \rightarrow \mathbf{M}_{\mathbf{g},r}$ assigns to Γ the distinguished component of Ω/Γ . Thus, there exists another projection map $\pi_2 : \mathbf{S}'_{\mathbf{g},r} \rightarrow \mathbf{M}_{\mathbf{g},r}$ that assigns to Γ the other Riemann surface. The product of π_1 and π_2 defines a holomorphic étale covering map

$$\mathbf{S}'_{\mathbf{g},r} \rightarrow \mathbf{M}_{\mathbf{g},r} \times \mathbf{M}_{\mathbf{g},r}$$

which one may think of as follows: The fiber over a point $([\mathbf{X}], [\mathbf{Y}]) \in \mathbf{M}_{\mathbf{g},r} \times \mathbf{M}_{\mathbf{g},r}$ is the set of homotopy classes of orientation-reversing homeomorphisms between the topological spaces underlying \mathbf{X} and \mathbf{Y} . Thus, relative to this description, $\mathrm{Im}(s_{\mathbf{H}}) \subseteq \mathbf{S}'_{\mathbf{g},r}$ consists of pairs $([\mathbf{X}], [\mathbf{X}^c]) \in \mathbf{M}_{\mathbf{g},r} \times \mathbf{M}_{\mathbf{g},r}$ (where the superscripted “c” denotes the complex conjugate Riemann surface), equipped with the tautological identification of the underlying topological spaces of \mathbf{X} and \mathbf{X}^c .

Now, let $\widetilde{\mathbf{M}}_{\mathbf{g},r} \rightarrow \mathbf{M}_{\mathbf{g},r}$ be the universal covering space of $\mathbf{M}_{\mathbf{g},r}$, and let $\widetilde{\mathbf{S}}'_{\mathbf{g},r} \rightarrow \widetilde{\mathbf{M}}_{\mathbf{g},r}$ be the pull-back of $\mathbf{S}'_{\mathbf{g},r} \rightarrow \mathbf{M}_{\mathbf{g},r}$ to $\widetilde{\mathbf{M}}_{\mathbf{g},r}$. Since it is known that the fibers of $\mathbf{S}'_{\mathbf{g},r} \rightarrow \mathbf{M}_{\mathbf{g},r}$ are simply connected, it follows that $\widetilde{\mathbf{S}}'_{\mathbf{g},r}$ is the universal covering space of $\mathbf{S}'_{\mathbf{g},r}$. Let Π be a standard copy of the fundamental group of a Riemann surface of type (g, r) . Then $\widetilde{\mathbf{S}}'_{\mathbf{g},r}$ can be described as the space of (isomorphism classes of) pairs of data (ρ, R) , where $\rho : \Pi \rightarrow \mathrm{PGL}_2(\mathbf{C})$ is an injective homomorphism whose

image is a quasi-Fuchsian group Γ of the sort described above, and R is a distinguished component of Ω/Γ .

Now we would like to define an anti-holomorphic involution

$$\Phi_{\tilde{S}} : \tilde{S}'_{g,r} \rightarrow \tilde{S}'_{g,r}$$

which we would like to think of as a *Frobenius action* on $\tilde{S}'_{g,r}$, as follows: The first datum of $\Phi_{\tilde{S}}(\rho, R)$ we take to be the representation

$$\bar{\rho}$$

(where the bar denotes complex conjugation). Let $\phi : \mathbf{P}^1(\mathbf{C}) \rightarrow \mathbf{P}^1(\mathbf{C})$ be the map that takes z to \bar{z} . Let Γ_{ϕ} be the image of $\bar{\rho}$ in $\mathrm{PGL}_2(\mathbf{C})$. Then the domain of discontinuity Ω_{ϕ} of Γ_{ϕ} is easily seen to be equal to $\phi(\Omega)$ (where Ω is the domain of discontinuity of $\Gamma \stackrel{\text{def}}{=} \mathrm{Im}(\rho)$). We take the distinguished component of Ω_{ϕ} to be the image under ϕ of the *nondistinguished* component of Ω . This completes the definition of $\Phi_{\tilde{S}}$. Moreover, it is clear from this definition that $\Phi_{\tilde{S}}$ is anti-holomorphic and of period two.

Next, let us observe that $\Phi_{\tilde{S}}$ descends to an anti-holomorphic involution Φ_S on $S'_{g,r}$. Moreover, $\pi_1 \circ \Phi_S$ is the composite of π_2 with complex conjugation on $M_{g,r}$. Note that this shows that Φ_S already essentially contains the data of the Bers uniformization: Indeed, the Bers uniformization, which is an anti-holomorphic map

$$S'_X \rightarrow M_{g,r}$$

is simply the result of composing the natural inclusion $S'_X \subseteq S'_{g,r}$ with Φ_S and then applying the natural projection $\pi_1 : S'_{g,r} \rightarrow M_{g,r}$. Put another way,

The Bers uniformization is, essentially, a Frobenius action. In particular, the Weil-Petersson metric, which is just the derivative (cf. [Mzk1], Introduction, §1) of the Bers uniformization, may be thought of as being essentially the Frobenius action on the tangent space of $M_{g,r}$.

Indeed, the point of view that the Weil-Petersson metric is essentially just a Frobenius action on the tangent space of $M_{g,r}$ is already inherent in Lemma 2.2 of [Mzk1], Introduction.

Now let $\tilde{s}_H : \tilde{M}_{g,r} \rightarrow \tilde{S}_{g,r}$ be the pull-back to $\tilde{M}_{g,r}$ of the real analytic section s_H . Then we claim that

The fixed point set of $\Phi_{\tilde{\mathbf{S}}}$ in $\tilde{\mathbf{S}}'_{\mathbf{g},\mathbf{r}} \subseteq \tilde{\mathbf{S}}_{\mathbf{g},\mathbf{r}}$ is precisely $\text{Im}(\tilde{s}_{\mathbf{H}})$.

Indeed, it is clear from the definition that $\Phi_{\tilde{\mathbf{S}}}$ fixes $\text{Im}(\tilde{s}_{\mathbf{H}})$. Now suppose that (ρ, R) is fixed by $\Phi_{\tilde{\mathbf{S}}}$. Then it follows that there exists a matrix $T \in \text{GL}_2(\mathbf{C})$ such that $\bar{\rho} = T\rho T^{-1}$. Since $H^1(\text{Gal}(\mathbf{C}/\mathbf{R}), \text{PGL}_2)$ has two elements, it is easy to work out that this implies that (up to isomorphism) Γ is contained in either $\text{PGL}_2(\mathbf{R})$ or a maximal compact subgroup of $\text{PGL}_2(\mathbf{C})$. But since Γ is discrete and infinite, the latter case is impossible. Thus, we conclude that Γ is Fuchsian. This completes the proof of the above claim. Thus, this claim is a slightly stronger version of the statement that the canonical indigenous bundle is the unique integral Frobenius invariant indigenous bundle.

In short, one can summarize complex Teichmüller theory in the following form:

One starts with the algebraic Schwarz torsor $\mathcal{S}_{g,r} \rightarrow \mathcal{M}_{g,r}$. The Fuchsian and Bers uniformizations are completely determined and can be recovered from the canonical real analytic section $s_{\mathbf{H}} : \mathcal{M}_{g,r}(\mathbf{C}) \rightarrow \mathcal{S}_{g,r}(\mathbf{C})$. Moreover, this section may be constructed as follows: At the infinite prime, the complex analytic Schwarz torsor $\mathbf{S}_{\mathbf{g},\mathbf{r}} \rightarrow \mathbf{M}_{\mathbf{g},\mathbf{r}}$ has a natural “integral structure” $\mathbf{S}'_{\mathbf{g},\mathbf{r}} \subseteq \mathbf{S}_{\mathbf{g},\mathbf{r}}$. Moreover, over the universal covering space $\tilde{\mathbf{S}}'_{\mathbf{g},\mathbf{r}}$ of $\mathbf{S}'_{\mathbf{g},\mathbf{r}}$, there is a natural Frobenius action

$$\Phi_{\tilde{\mathbf{S}}} : \tilde{\mathbf{S}}'_{\mathbf{g},\mathbf{r}} \rightarrow \tilde{\mathbf{S}}'_{\mathbf{g},\mathbf{r}}$$

Then the canonical section $s_{\mathbf{H}}$ is precisely that section whose pull-back $\tilde{s}_{\mathbf{H}}$ to $\tilde{\mathbf{M}}_{\mathbf{g},\mathbf{r}}$ has image in $\tilde{\mathbf{S}}_{\mathbf{g},\mathbf{r}}$ equal to the invariant set of the Frobenius morphism $\Phi_{\tilde{\mathbf{S}}}$.

Put another way,

One may regard complex Teichmüller theory as the study of the canonical Frobenius action $\Phi_{\tilde{\mathbf{S}}}$ on $\tilde{\mathbf{S}}'_{\mathbf{g},\mathbf{r}}$.

We shall see that in this form, Teichmüller theory carries over very nicely to finite primes.

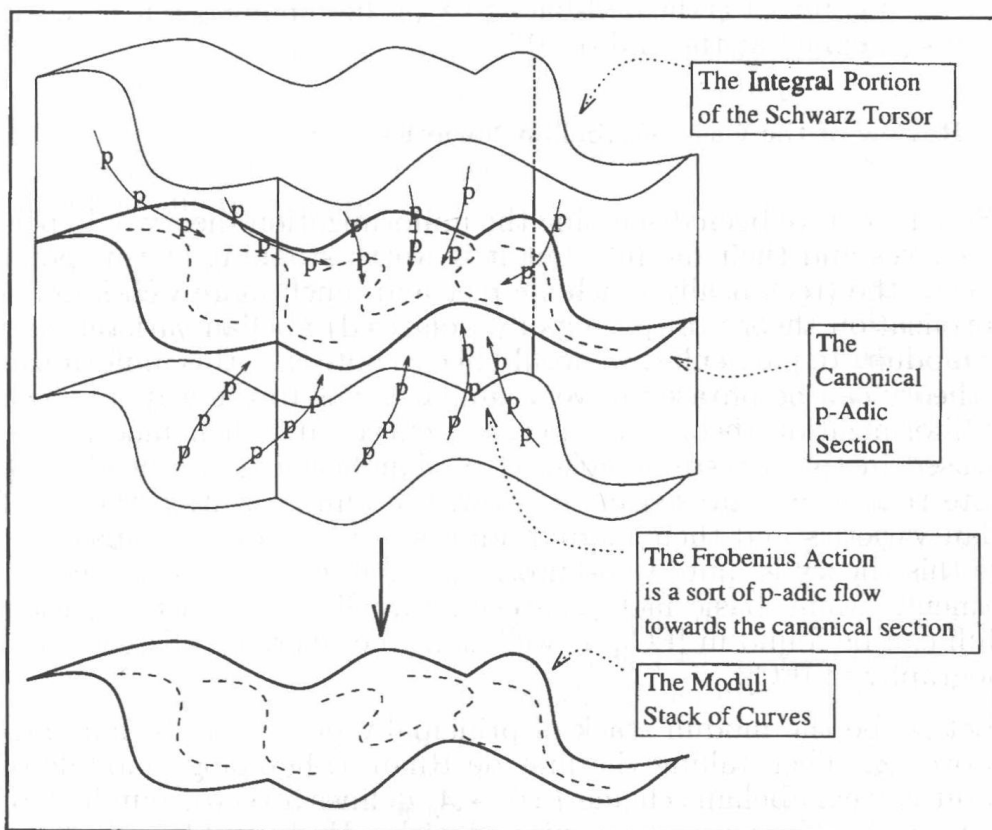


Fig. 3: The Canonical Frobenius Action on the Integral Portion of the Schwarz Torsor

Finally, let us note that on an individual curve X , one has a “Frobenius action” that lies at the center of the uniformization theory of the curve: Namely, let $T \rightarrow X$ be the torsor discussed at the end of §0.5; let $\mathbf{T} \rightarrow \mathbf{X}$ be the corresponding analytic object. Observe that since the monodromy of the canonical indigenous bundle $(\mathbf{P} \rightarrow \mathbf{X}, \nabla_{\mathbf{P}})$ is real, there exists a natural subbundle $\mathbf{P}_{\mathbf{R}} \subseteq \mathbf{P}$ of real projective spaces over \mathbf{X} . Then $\mathbf{P} - \mathbf{P}_{\mathbf{R}}$ consists of two connected components, one of which is contained in $\mathbf{T} \subseteq \mathbf{P}$. We shall denote this connected component by $\mathbf{T}' \subseteq \mathbf{T}$, and think of it as defining an integral structure on \mathbf{T} . Then it follows easily from the construction of the canonical indigenous bundle that the pull-back $\tilde{\mathbf{T}}' \rightarrow \tilde{\mathbf{X}}$ of $\mathbf{T}' \rightarrow \mathbf{X}$ is naturally isomorphic to the projection $\mathbf{H} \times \mathbf{H}^c \rightarrow \mathbf{H}$ (where \mathbf{H}^c denotes the *lower* half plane) to the first factor. Moreover, there is a natural Frobenius action on $\mathbf{H} \times \mathbf{H}^c$ (i.e., the product of the upper and lower half planes of \mathbf{C}) given by sending (z_1, z_2) to (\bar{z}_2, \bar{z}_1) . This defines a Frobenius action

$$\Phi_{\tilde{\mathbf{T}}} : \tilde{\mathbf{T}}' \rightarrow \tilde{\mathbf{T}}'$$

whose fixed point set is the pull-back to $\tilde{\mathbf{X}}$ of the canonical real analytic section $s_{\mathbf{X}}$ defined at the end of §0.5.

§0.7. Review of the Case of Abelian Varieties

So far we have been discussing the uniformization theory of hyperbolic curves and their moduli, but it is worth stopping at this point to review the (technically much simpler and much more well-known) uniformization theory of (principally polarized) abelian varieties and their moduli. In particular, we would like to note that this uniformization theory can be phrased in very similar terms to the way in which the uniformization theory of hyperbolic curves and their moduli was rephrased in the preceding subsection. Furthermore, we would like to note that when phrased in this way, the uniformization theory of abelian varieties and their moduli admits a natural p -adic analogue. Since this theory is quite well-known, we shall give a rather succinct treatment. Many basic facts concerning abelian varieties and their moduli can be found in [FC], as well as in the references listed in the bibliography of [FC].

Let \mathcal{A}_g be the moduli stack of principally polarized abelian varieties over \mathbf{Z} . Then taking the first de Rham cohomology module of the tautological abelian scheme $f : \mathcal{G} \rightarrow \mathcal{A}_g$ defines a vector bundle \mathcal{E} of rank $2g$ on \mathcal{A}_g . Moreover, \mathcal{E} is equipped with a Hodge subbundle $\mathcal{F} \subseteq \mathcal{E}$ such that \mathcal{F} is naturally isomorphic to $f_* \Omega_{\mathcal{G}/\mathcal{A}_g}$. The principal polarization on \mathcal{G} defines a natural isomorphism between \mathcal{F} and the dual of \mathcal{E}/\mathcal{F} ; we shall thus identify \mathcal{F} with the dual of \mathcal{E}/\mathcal{F} in the sequel. Thus, the splittings of $\mathcal{F} \subseteq \mathcal{E}$ form a torsor $\mathcal{T}_1 \rightarrow \mathcal{A}_g$ over $\mathcal{F} \otimes_{\mathcal{O}_{\mathcal{A}_g}} \mathcal{F}$. Pushing forward this torsor by the natural quotient $\mathcal{F} \otimes_{\mathcal{O}_{\mathcal{A}_g}} \mathcal{F} \rightarrow \mathbf{S}^2 \mathcal{F}$ (to the second symmetric power of \mathcal{F}) gives a torsor $\mathcal{T} \rightarrow \mathcal{A}_g$. The Kodaira-Spencer morphism given by differentiating the filtration $\mathcal{F} \subseteq \mathcal{E}$ by means of the Gauss-Manin connection on \mathcal{E} defines a natural isomorphism of $\mathbf{S}^2 \mathcal{F}$ with $\Omega_{\mathcal{A}_g/\mathbf{Z}}$. Thus, we may regard $\mathcal{T} \rightarrow \mathcal{A}_g$ as a torsor over $\Omega_{\mathcal{A}_g/\mathbf{Z}}$.

Now we would like to look at things at the infinite prime. We shall denote by means of boldface letters the complex analytic objects defined by the various algebraic objects over \mathbf{Z} discussed in the preceding paragraph. First, let us observe that one has a “real structure” on the de Rham cohomology module \mathbf{E} , given by the de Rham isomorphism which exhibits \mathbf{E} as the result of tensoring (over \mathbf{R}) the relative real singular cohomology of the tautological abelian analytic stack $\mathbf{G} \rightarrow \mathbf{A}_g$ with $\mathcal{O}_{\mathbf{A}_g}$. This real structure allows us to regard \mathbf{E} – *in the real analytic category* – as the direct sum of \mathbf{F} and its complex conjugate $\overline{\mathbf{F}}$. There are several things that one can do with this direct sum decomposition $\mathbf{E} = \mathbf{F} \oplus \overline{\mathbf{F}}$. First, note that by projection, we get an isomorphism (in the real analytic category) $\overline{\mathbf{F}} \rightarrow \mathbf{E}/\mathbf{F} = (\mathbf{F})^\vee$, i.e., a Hermitian form on

\mathbf{F} . This form in fact defines a real analytic metric on $\mathbf{F} = f_* \Omega_{\mathbf{G}/\mathbf{A}_g}$ – which we will henceforth refer to as the *standard metric on \mathbf{F}* – and which will be of use to us below. Moreover, by functoriality, it also defines a real analytic metric on $\mathbf{S}^2 \mathbf{F} = \Omega_{\mathbf{A}_g/\mathbf{C}}$. It is easy to see that this metric is precisely the metric $\mu_{\mathbf{A}_g}$ arising from the Siegel upper half plane uniformization.

The next thing that this real analytic direct sum decomposition $\mathbf{E} = \mathbf{F} \oplus \overline{\mathbf{F}}$ is good for is that it defines (by the definition of \mathcal{T}_1) a real analytic section of $\mathcal{T}_1 \rightarrow \mathcal{A}_g$, and, in particular, a real analytic section s_∞ of $\mathcal{T} \rightarrow \mathcal{A}_g$. It is an elementary exercise in linear algebra (cf. Theorem 2.3 of [Mzk1], Introduction, §2) to verify that $\bar{\partial}s_\infty$ of this real analytic section is equal to the metric $\mu_{\mathbf{A}_g}$ arising from the Siegel upper half plane uniformization. (Indeed, one can already see that $\bar{\partial}s_\infty$ and $\mu_{\mathbf{A}_g}$ differ at most by a nonzero constant multiple from the fact that the pull-back of s_∞ , and hence of $\bar{\partial}s_\infty$, to $\tilde{\mathbf{A}}_g$ (the universal covering space of \mathbf{A}_g) is (by its definition) manifestly invariant under the natural action of $\mathrm{Sp}_{2g}(\mathbf{R})$ on $\tilde{\mathbf{A}}_g$.) Incidentally, this also shows that the torsor $\mathcal{T} \rightarrow \mathcal{A}_g$ is the torsor associated to (some positive tensor power of) the ample line bundle $\det(\mathcal{F})$. Thus, in summary,

Just as in the case of the Schwarz torsor on $\mathcal{M}_{g,r}$, we have an algebraic torsor $\mathcal{T} \rightarrow \mathcal{A}_g$ (which arises as the Hodge-theoretic first Chern class of an ample line bundle) equipped with a canonical real analytic section s_∞ that embodies the uniformization theory of \mathbf{A}_g .

Moreover, just as in the case of moduli of hyperbolic curves, there are natural integral structures and Frobenius actions defined as follows: Indeed, to see this, first observe that $\mathcal{T} \rightarrow \mathcal{A}_g$ can also be regarded as parametrizing subbundles $\mathcal{F}' \subseteq \mathcal{E}$ such that $\mathcal{E} = \mathcal{F} \oplus \mathcal{F}'$ and, moreover, the Riemann form on \mathcal{E} (arising from the principal polarization) is trivial on \mathcal{F}' . Then one obtains an integral structure \mathbf{T}' on the corresponding analytic torsor $\mathbf{T} \rightarrow \mathbf{A}_g$ by taking \mathbf{T}' to be the set of $\mathcal{F}' \subseteq \mathcal{E}$ such that \mathcal{E} , equipped with (a.) its lattice (i.e., \mathbf{Z} -) structure, (b.) its Riemann form, and (c.) the complex conjugate subspace $\overline{\mathcal{F}'} \subseteq \mathcal{E}$ to \mathcal{F}' , defines a principally polarized abelian variety. Then the Frobenius action $\Phi_{\mathbf{T}}$ on \mathbf{T}' is the map that sends $(\mathcal{F} \subseteq \mathcal{E}, \mathcal{F}' \subseteq \mathcal{E})$ to $(\overline{\mathcal{F}'} \subseteq \mathcal{E}, \overline{\mathcal{F}} \subseteq \mathcal{E})$. One checks easily that $\Phi_{\mathbf{T}}$ is an anti-holomorphic involution; that (just as in the case of moduli of hyperbolic curves) it essentially defines the Siegel upper half plane uniformization of $\tilde{\mathbf{A}}_g$; and that the analogously defined Frobenius action on the universal covering $\tilde{\mathbf{T}}'$ of \mathbf{T}' has as its fixed point set precisely the image of the (pull-back to $\tilde{\mathbf{A}}_g$ of the) canonical real analytic section s_∞ defined above.

Just as in the case of hyperbolic curves, there is a similar construction that one can carry out over individual (principally polarized)

abelian varieties. Indeed, let A be a principally polarized abelian variety over \mathbf{C} . Then one has a *universal vector extension* $0 \rightarrow V \rightarrow U \rightarrow A \rightarrow 0$ over A , where V is the affine group scheme over \mathbf{C} corresponding to the complex vector space $H^0(A^t, \Omega_{A^t/\mathbf{C}})$, and A^t is the dual abelian variety to A . Since A is principally polarized, we have an isomorphism $A^t \cong A$, so we may think of V as corresponding to $H^0(A, \Omega_{A/\mathbf{C}})$. Then the extension U defines an $\Omega_{A/\mathbf{C}}$ -torsor $T_A \rightarrow A$ over A which is the Hodge-theoretic first Chern class of the line bundle defining the polarization $A \cong A^t$. Now let $A' \subseteq U$ be the closure (in the complex analytic topology) of the set of torsion points of U . Then it is easy to show that A' is a real torus that maps isomorphically via the projection $U \rightarrow A$ onto the real torus underlying A . In other words, A' defines a real analytic section s_A of the torsor $T_A \rightarrow A$. As usual, it is a matter of elementary linear algebra to show that $\bar{\partial}s_A$ is the metric on $\Omega_{A/\mathbf{C}}$ induced by the standard metric on $H^0(A, \Omega_{A/\mathbf{C}})$ (defined two paragraphs above). Thus,

Just as in the case of hyperbolic curves, we have a torsor $T_A \rightarrow A$ (which arises as the Hodge-theoretic first Chern class of an ample line bundle) equipped with a canonical real analytic section s_A that embodies the uniformization theory of A .

Moreover, one can define integral structures and Frobenius actions as follows: First, of all we take the “integral structure” on $\mathbf{T}_A \rightarrow \mathbf{A}$ to be \mathbf{T}_A itself. (This sort of “degenerate” integral structure arises from the fact that unlike the other cases considered so far, \mathbf{A} is not a hyperbolic space.) Then \mathbf{T}_A is naturally uniformized by the affine space defined by the singular homology module $H_1^{\text{sing}}(\mathbf{A}, \mathbf{C})$. This homology module comes equipped with a natural real structure given by $H_1^{\text{sing}}(\mathbf{A}, \mathbf{R})$, hence a natural complex conjugation morphism. This complex conjugation morphism thus defines a natural Frobenius action on the universal covering space $\tilde{\mathbf{T}}_A$ of \mathbf{T}_A . Moreover, this Frobenius action has all the usual properties (i.e., essentially defines the uniformization morphism, its fixed point set is the image of the pull-back of s_A to $\tilde{\mathbf{A}}$, etc.).

Now, before proceeding, let us observe that, at least in the case of (principally polarized) abelian varieties and their moduli, there is a well-known p -adic analogue of the theory just discussed at the infinite prime. Indeed, let p be a prime number, and let $\mathcal{A}_g^{\text{ord}}$ be the p -adic formal stack which is the p -adic completion of \mathcal{A}_g at the ordinary locus of $\mathcal{A}_g \otimes \mathbf{F}_p$. Then we have a morphism $\Phi_A : \mathcal{A}_g^{\text{ord}} \rightarrow \mathcal{A}_g^{\text{ord}}$ given by mapping an abelian scheme G (with ordinary reduction modulo p) to the abelian scheme which is the quotient of G by the multiplicative portion of the kernel of p on G . (Here, by “multiplicative portion,” we mean the maximal subgroup scheme which is étale locally isomorphic to a product of μ_p ’s, where μ_p is the kernel of p on the multiplicative

group scheme \mathbf{G}_m .) Then $\Phi_A \otimes \mathbf{F}_p$ is equal to the Frobenius morphism, i.e., Φ_A is a *Frobenius lifting*. Moreover, we have a natural Frobenius action $\Phi_{\mathcal{E}} : \Phi_A^* \mathcal{E}|_{\mathcal{A}_g^{\text{ord}}} \rightarrow \mathcal{E}|_{\mathcal{A}_g^{\text{ord}}}$ on $\mathcal{E}|_{\mathcal{A}_g^{\text{ord}}}$ which is divisible by p on $\Phi_A^* \mathcal{F}|_{\mathcal{A}_g^{\text{ord}}}$. This Frobenius action induces a Frobenius action on $\mathcal{T}_1|_{\mathcal{A}_g^{\text{ord}}}$ and $\mathcal{T}|_{\mathcal{A}_g^{\text{ord}}}$, i.e., (in the latter case) a morphism $\Phi_{\mathcal{T}} : \mathcal{T}|_{\mathcal{A}_g^{\text{ord}}} \rightarrow \mathcal{T}|_{\mathcal{A}_g^{\text{ord}}}$ lying over Φ_A . Moreover, the torsor $\mathcal{T}|_{\mathcal{A}_g^{\text{ord}}} \rightarrow \mathcal{A}_g^{\text{ord}}$ admits a unique section $s_p : \mathcal{A}_g^{\text{ord}} \rightarrow \mathcal{T}|_{\mathcal{A}_g^{\text{ord}}}$ which is $\Phi_{\mathcal{T}}$ -invariant. Put another way, this invariant section s_p corresponds to the subspace of $\mathcal{E}|_{\mathcal{A}_g^{\text{ord}}}$ that is usually referred to as the *unit Frobenius subspace*. Thus, we see that

The theory of ordinary abelian varieties gives us a canonical p -adic analytic trivialization of the pull-back of the torsor $\mathcal{T} \rightarrow \mathcal{A}_g$ to $\mathcal{A}_g^{\text{ord}}$.

We pause here to remark that the analogy between s_{∞} and s_p is studied in great detail in the one-dimensional case in [Katz] (see especially, §1.3.6, 5.7.7).

Next, note that the definition of Φ_A is such that we get a natural endomorphism $\Phi_{\mathcal{G}} : \mathcal{G}|_{\mathcal{A}_g^{\text{ord}}} \rightarrow \mathcal{G}|_{\mathcal{A}_g^{\text{ord}}}$ (lying over Φ_A) of the pull-back of the tautological abelian scheme to $\mathcal{A}_g^{\text{ord}}$. This morphism $\Phi_{\mathcal{G}}$ is a lifting of Frobenius (i.e., modulo p , it is equal to the Frobenius morphism). Moreover, if $\mathcal{U} \rightarrow \mathcal{G}$ is the universal vector extension of \mathcal{G} , then it follows from the universality property of the universal vector extension that one gets a natural morphism $\Phi_{\mathcal{U}} : \mathcal{U}|_{\mathcal{A}_g^{\text{ord}}} \rightarrow \mathcal{U}|_{\mathcal{A}_g^{\text{ord}}}$ lying over $\Phi_{\mathcal{G}}$. It is not difficult to see that $\mathcal{U}|_{\mathcal{A}_g^{\text{ord}}} \rightarrow \mathcal{G}|_{\mathcal{A}_g^{\text{ord}}}$ admits a unique $\Phi_{\mathcal{U}}\text{-}\Phi_{\mathcal{G}}$ -equivariant section in a formal neighborhood of the zero section of $\mathcal{G}|_{\mathcal{A}_g^{\text{ord}}} \rightarrow \mathcal{A}_g^{\text{ord}}$. Thus, we see that

In the case of individual ordinary (principally polarized) abelian varieties, we again obtain a canonical p -adic analytic trivialization of the pull-back of the torsor over $\mathcal{G}|_{\mathcal{A}_g^{\text{ord}}}$ which is the Hodge-theoretic first Chern class of the polarization.

In the following subsection, we shall observe that this phenomenon of analogous complex and p -adic theories is by no means limited to the case of abelian varieties and their moduli.

§0.8. Arithmetic Frobenius Venues

At this point, it is useful to take a step back and look at what is going on here in a more general context. Namely, consider the following situation: We begin with a smooth scheme M defined over \mathbf{Q} (hence defined over most \mathbf{Z}_p). Suppose further that this scheme

M is equipped with an ample line bundle \mathcal{L} (defined over \mathbf{Q}). Then the torsor of connections on \mathcal{L} may be realized geometrically as a morphism $T \rightarrow M$. In general, $T \rightarrow M$ will not admit any algebraic sections, or even any holomorphic sections (after base-change to \mathbf{C}). Next, let us suppose that the corresponding analytic torsor $\mathbf{T} \rightarrow \mathbf{M}$ is equipped with an “integral structure” $\mathbf{T}' \subseteq \mathbf{T}$. That is, $\mathbf{T}' \subseteq \mathbf{T}$ is an open subset such that the restriction of \mathbf{T}' to each fiber of $\mathbf{T} \rightarrow \mathbf{M}$ is a nonempty open, simply connected subset of that fiber. Next, let us assume that we are given a *Frobenius action*, i.e., an anti-holomorphic involution

$$\Phi_{\tilde{\mathbf{T}}} : \tilde{\mathbf{T}}' \rightarrow \tilde{\mathbf{T}}'$$

on the pull-back of $\mathbf{T}' \rightarrow \mathbf{M}$ to the universal covering space $\tilde{\mathbf{M}}$ of \mathbf{M} that descends to a Frobenius action $\Phi_{\mathbf{T}}$ on \mathbf{T}' . Moreover, let us assume that the fixed point locus of $\Phi_{\tilde{\mathbf{T}}}$ is the image of the pull-back to $\tilde{\mathbf{M}}$ of a *real analytic section* $s_{\infty} : \mathbf{M} \rightarrow \mathbf{T}$ whose derivative $\bar{\partial}s_{\infty}$ defines a (real analytic) Kähler metric on \mathbf{M} .

We shall refer to such a triple of data $(M, \mathcal{L}, \Phi_{\tilde{\mathbf{T}}} : \tilde{\mathbf{T}}' \rightarrow \tilde{\mathbf{T}}')$ as an arithmetic Kähler venue.

Note that just as in all the cases we have seen so far – i.e., moduli of hyperbolic curves and abelian varieties, hyperbolic curves, abelian varieties – it is a *tautology* (cf. Theorem 2.3 of [Mzk1], Introduction; Theorem 2.6 of §2 of the Appendix to the present book) that

The Kähler metric $\bar{\partial}s_{\infty}$ is precisely the Hermitian form induced on $\Omega_{\mathbf{M}}$ by composing the inclusion $\Omega_{\mathbf{M}} \hookrightarrow \Omega_{\mathbf{T}}|_{\text{Im}(s_{\infty})}$ (induced by the projection $\mathbf{T} \rightarrow \mathbf{M}$) with $d\Phi_{\mathbf{T}}$, followed by the projection $\Omega_{\mathbf{T}}|_{\text{Im}(s_{\infty})} \rightarrow \Omega_{\mathbf{M}}^{\vee}$ (derived from the fact that $\mathbf{T} \rightarrow \mathbf{M}$ is an $\Omega_{\mathbf{M}}$ -torsor).

Indeed, the subspace of $\Omega_{\mathbf{T}}$ (at a point in the image of s_{∞}) corresponding to the image of s_{∞} is precisely the “real structure” on $\Omega_{\mathbf{T}}$ given by taking the invariants of the Frobenius action $d\Phi_{\mathbf{T}}$ on $\Omega_{\mathbf{T}}$. Now it is a matter of elementary linear algebra that when one has a real vector space V whose complexification $V_{\mathbf{C}} \stackrel{\text{def}}{=} V \otimes_{\mathbf{R}} \mathbf{C}$ is equipped with a filtration $F \subseteq V_{\mathbf{C}}$ (i.e., F is a \mathbf{C} -subspace) and a \mathbf{C} -isomorphism $F \cong (V_{\mathbf{C}}/F)^{\vee}$ such that the composite $V \subseteq V_{\mathbf{C}} \rightarrow Q \stackrel{\text{def}}{=} V_{\mathbf{C}}/F$ is bijective, then $\bar{\partial}$ of the “real analytic section” of $V_{\mathbf{C}} \rightarrow Q$ given by $V \subseteq V_{\mathbf{C}}$ is given by the recipe described above. This completes the justification of the “tautology” just cited. That is to say,

The Kähler metric $\bar{\partial}s_{\infty}$ is essentially the derivative of the Frobenius morphism $\Phi_{\mathbf{T}}$.

The first goal of this subsection is to point out that natural arithmetic Kähler venues exist in great abundance. Indeed, so far, we have seen examples where M is:

- (1) the moduli stack of r -pointed smooth curves of genus g (where we assume $2g - 2 + r > 0$);
- (2) an arbitrary hyperbolic curve;
- (3) the moduli stack of principally polarized abelian varieties;
- (4) an arbitrary principally polarized abelian variety.

A less well-known example (see the Appendix to this book, especially §2, Theorem 2.6, for more details) is the case where M is:

- (5) the moduli space of stable bundles on an algebraic curve.

What is significant is that in all but the first two examples mentioned above,

There are well-known p -adic analogues of $\Phi_{\tilde{T}}$ and s_{∞} .

Indeed, this p -adic theory was reviewed in the preceding subsection in the case of abelian varieties and their moduli; moreover, in the least well-known of these examples – the moduli space of stable bundles – the p -adic analogue is discussed in detail in §3 (see especially Corollary 3.13) of the Appendix to this book. Thus, it is natural to expect that in the first two examples as well – i.e., hyperbolic curves and their moduli – there should exist canonical p -adic analytic trivializations of the torsor $T \rightarrow M$.

Next, let us note that in all the p -adic examples discussed so far, the canonical p -adic analytic trivializations s_p were obtained as follows:

- (a.) One first constructs some sort of p -adic, typically “ordinary,” version – let us call it M_p – of M .
- (b.) One defines a canonical Frobenius action on the pull-back of the torsor $T \rightarrow M$ to M_p .
- (c.) Then one obtains the canonical trivialization as the *unique trivialization of the torsor $T|_{M_p} \rightarrow M_p$ fixed by the Frobenius action*.

The point here is that

The above recipe in the p -adic case is completely formally analogous to the recipe by which the s_∞ 's were obtained, i.e., by looking at fixed point sets of some canonical Frobenius action $\Phi_{\tilde{T}}$.

In fact, the condition (which, in the complex case, tends to be automatically satisfied) that the fixed point set of $\Phi_{\tilde{T}}$ be transversal to the projection $\tilde{T}' \rightarrow \tilde{M}$ may even be regarded as a condition of *ordinariness* in the following sense: The definition of “ordinariness” in p -adic contexts (e.g., abelian varieties; stable bundles as discussed in §3 of the Appendix to this book; the theory of [Mzk1] and the present book) amounts to the condition that the image of the Frobenius action on the tangent space to T in characteristic p be transversal to the filtration on that tangent space arising from the projection $T \rightarrow M$.

In fact, in the case of hyperbolic curves and their moduli considered in [Mzk1] and the present work, we shall obtain canonical p -adic trivializations by means of precisely this recipe of looking at Frobenius invariants.

Finally, let us observe that

There is a certain analogy between the Frobenius action on M_p that one typically gets in the p -adic case and the Kähler metric that one gets in the complex case.

Indeed, both the Frobenius action on M_p and the Kähler metric are additional structures on M that are obtained by essentially restricting the Frobenius action on T to M . Moreover, just as the Kähler metric gives rise (via formal integration) to canonical coordinates on M , the Frobenius action on M_p typically (in particular, in all the examples listed above) gives rise (via a sort of formal integration – cf. [Mzk1], Chapter III, §1; Chapters VIII and IX of the present work) to p -adic canonical coordinates on M . For instance, in the case of moduli of principally polarized varieties, these coordinates are the Serre-Tate coordinates; in the case of abelian varieties, the canonical coordinates are essentially those obtained from the exponential map; in the case of moduli of stable bundles, the canonical p -adic coordinates are discussed in §3 of the Appendix to this book.

Thus, in summary, it is the opinion of the author that

There is a strong analogy between the “arithmetic Kähler venues” that arise in the complex case and the corresponding situations that arise p -adically.

It is for this reason that we feel that it is natural to refer to “arithmetic Kähler venues” as *arithmetic Frobenius venues at the infinite prime*, and to

refer to their p -adic counterparts as *arithmetic Frobenius venues at finite primes*.

§0.9. The Classical Ordinary Theory

As stated earlier, the purpose of the present work is to study all integral Frobenius invariant indigenous bundles. On the other hand, in [Mzk1], a very important special type of Frobenius invariant indigenous bundle was constructed. This type of bundle will henceforth be referred to as *classical ordinary*. (Such bundles were called “ordinary” in [Mzk1]. Here we use the term “classical ordinary” to refer to objects called “ordinary” in [Mzk1] in order to avoid confusion with the more general notions of ordinariness discussed in the present work.) Before discussing the theory of the present book (which is the goal of §1), it is thus natural to review the classical ordinary theory. In this subsection, we let p be an odd prime.

If one is to construct p -adic Frobenius invariant indigenous bundles for arbitrary hyperbolic curves, the first order of business is to make precise the notion of Frobenius invariance that one is to use. For this, it is useful to have a prototype. The prototype that gave rise to the classical ordinary theory is the following:

Let $\mathcal{M} \stackrel{\text{def}}{=} (\mathcal{M}_{1,0})_{\mathbf{Z}_p}$ be the moduli stack of elliptic curves over \mathbf{Z}_p . Let $\mathcal{G} \rightarrow \mathcal{M}$ be the universal elliptic curve. Let \mathcal{E} be its first de Rham cohomology module. Thus, \mathcal{E} is a rank two vector bundle on \mathcal{M} , equipped with a Hodge subbundle $\mathcal{F} \subseteq \mathcal{E}$, and a connection $\nabla_{\mathcal{E}}$ (i.e., the “Gauss-Manin connection”). Taking the projectivization of \mathcal{E} defines a \mathbf{P}^1 -bundle with connection $(P \rightarrow \mathcal{M}, \nabla_P)$, together with a Hodge section $\sigma : \mathcal{M} \rightarrow P$. It turns out that (the natural extension over the compactification of \mathcal{M} obtained by using log structures of) the bundle (P, ∇_P) is an *indigenous bundle* on \mathcal{M} . In particular, (P, ∇_P) defines a crystal in \mathbf{P}^1 -bundles on $\text{Crys}(\mathcal{M} \otimes \mathbf{F}_p/\mathbf{Z}_p)$. Thus, one can form the pull-back $\Phi^*(P, \nabla_P)$ via the Frobenius morphism of this crystal. If one then adjusts the integral structure of $\Phi^*(P, \nabla_P)$ (cf. Definition 1.18 of Chapter VI of the present work; [Mzk1], Chapter III, Definition 2.4), one obtains the *renormalized Frobenius pull-back* $\mathbf{F}^*(P, \nabla_P)$. Then (P, ∇_P) is Frobenius invariant in the sense that $(P, \nabla_P) \cong \mathbf{F}^*(P, \nabla_P)$.

Thus, the basic idea behind [Mzk1] was to consider to what extent one could construct indigenous bundles on arbitrary hyperbolic curves that are equal to their own renormalized Frobenius pull-backs, i.e., satisfying

$$\mathbf{F}^*(P, \nabla_P) \cong (P, \nabla_P)$$

In particular, it is natural to try to consider moduli of indigenous bundles satisfying this condition. Since it is not at all obvious how to do this over \mathbf{Z}_p , a natural first step was to make the following *key* observation:

If (P, ∇_P) is an indigenous bundle over \mathbf{Z}_p preserved by \mathbf{F}^ , then the reduction modulo p of (P, ∇_P) has square nilpotent p -curvature.*

(The “ p -curvature” of an indigenous bundle in characteristic p is a natural invariant of such a bundle. We refer to [Mzk1], Chapter II, as well as §1 of Chapter II of the present book for more details.) Thus, if $(\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ is the stack of r -pointed stable curves of genus g in characteristic p , one can define the stack $\overline{\mathcal{N}}_{g,r}$ of such curves equipped with a “nilpotent” indigenous bundle. (Here, “nilpotent” means that its p -curvature is square nilpotent.) In the following, we shall often find it convenient to refer to pointed stable curves equipped with nilpotent indigenous bundles as *nilcurves*, for short. Thus, $\overline{\mathcal{N}}_{g,r}$ is the moduli stack of nilcurves. We would like to emphasize that

The above observation – which led to the notion of “nilcurves” – is the key technical breakthrough that led to the development of the “ p -adic Teichmüller theory” of [Mzk1] and the present work.

The first major result of [Mzk1] is the following (cf. Chapter II, Proposition 1.7; [Mzk1], Chapter II, Theorem 2.3):

Theorem 0.1. (Stack of Nilcurves) *The natural morphism $\overline{\mathcal{N}}_{g,r} \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ is a finite, flat, local complete intersection morphism of degree p^{3g-3+r} .*

In particular, up to “isogeny” (i.e., up to the fact that $p^{3g-3+r} \neq 1$), the stack of nilcurves $\overline{\mathcal{N}}_{g,r} \subseteq \overline{\mathcal{S}}_{g,r}$ defines a canonical section of the Schwarz torsor $\overline{\mathcal{S}}_{g,r} \rightarrow \overline{\mathcal{M}}_{g,r}$ in characteristic p .

Thus, relative to our discussion of complex Teichmüller theory – which we saw could be regarded as the study of a certain canonical real analytic section of the Schwarz torsor – it is natural that “ p -adic Teichmüller theory” should revolve around the study of $\overline{\mathcal{N}}_{g,r}$.

Although the structure of $\overline{\mathcal{N}}_{g,r}$ is now been much better understood, at the time of writing of [Mzk1] (Spring of 1994), it was not so well understood, and so it was natural to do the following: Let $\overline{\mathcal{N}}_{g,r}^{\text{ord}} \subseteq \overline{\mathcal{N}}_{g,r}$ be the open substack where $\overline{\mathcal{N}}_{g,r}$ is *étale* over $(\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$.

This open substack will be referred to as the *(classical) ordinary locus* of $\overline{\mathcal{N}}_{g,r}$. It is worth pausing here to note the following: The reason for the use of the term “ordinary” is that it is standard general practice to refer to as “ordinary” situations where Frobenius acts on a linear space equipped with a “Hodge subspace” in such a way that it acts with slope zero on a subspace of the same rank as the rank of the Hodge subspace. Thus, we use the term “ordinary” here because the Frobenius action on the cohomology of an ordinary nilcurve satisfies just such a condition. In other words, ordinary nilcurves are ordinary in their capacity as nilcurves. However, it is important to remember that:

The issue of whether or not a nilcurve is ordinary is entirely different from the issue of whether or not the Jacobian of the underlying curve is ordinary (in the usual sense). That is to say, there exist examples of ordinary nilcurves whose underlying curves have nonordinary Jacobians as well as examples of nonordinary nilcurves whose underlying curves have ordinary Jacobians.

Later, we shall comment further on the issue of the incompatibility of the theory of [Mzk1] with Serre-Tate theory relative to the operation of passing to the Jacobian.

At any rate, since $\overline{\mathcal{N}}_{g,r}^{\text{ord}}$ is étale over $(\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$, it lifts naturally to a p -adic formal stack \mathcal{N} which is étale over $(\overline{\mathcal{M}}_{g,r})_{\mathbf{Z}_p}$. Let $\mathcal{C} \rightarrow \mathcal{N}$ denote the tautological stable curve over \mathcal{N} . Then the main result (Theorem 0.1 of the Introduction of [Mzk1]) of the theory of [Mzk1] is the following:

Theorem 0.2. (Canonical Frobenius Lifting) *There exists a unique pair $(\Phi_{\mathcal{N}} : \mathcal{N} \rightarrow \mathcal{N}; (P, \nabla_P))$ satisfying the following:*

- (1) *The reduction modulo p of the morphism $\Phi_{\mathcal{N}}$ is the Frobenius morphism on \mathcal{N} , i.e., $\Phi_{\mathcal{N}}$ is a Frobenius lifting.*
- (2) *(P, ∇_P) is an indigenous bundle on \mathcal{C} such that the renormalized Frobenius pull-back of $\Phi_{\mathcal{N}}^*(P, \nabla_P)$ is isomorphic to (P, ∇_P) , i.e., (P, ∇_P) is Frobenius invariant with respect to $\Phi_{\mathcal{N}}$.*

Moreover, this pair also gives rise in a natural way to a Frobenius lifting $\Phi_{\mathcal{C}} : \mathcal{C}^{\text{ord}} \rightarrow \mathcal{C}^{\text{ord}}$ on a certain formal p -adic open substack \mathcal{C}^{ord} of \mathcal{C} (which will be referred to as the ordinary locus of \mathcal{C}).

Thus, this Theorem is a partial realization of the goal of constructing a canonical integral Frobenius invariant bundle on the universal stable curve.




Fig. 4: The Canonical Frobenius Action Underlying Theorem 0.2

Again, we observe that

This canonical Frobenius lifting $\Phi_{\mathcal{N}}$ is by no means compatible with the canonical Frobenius lifting $\Phi_{\mathcal{A}}$ discussed in §0.7 (relative to the operation of passing to the Jacobian).

At first glance, the reader may find this fact to be extremely disappointing and unnatural. In fact, however, when understood properly, this incompatibility is something which is to be expected. Indeed, relative to the analogy between Frobenius liftings and Kähler metrics discussed in §0.8 such a compatibility would be the p -adic analogue of a compatibility between the Weil-Petersson metric on $(\mathcal{M}_{g,r})_{\mathbb{C}}$ and the Siegel upper half plane metric on $(\mathcal{A}_g)_{\mathbb{C}}$. On the other hand, it is easy to see in the complex case that these two metrics are far from compatible. (Indeed, if they were compatible, then the Torelli map $(\mathcal{M}_g)_{\mathbb{C}} \rightarrow (\mathcal{A}_g)_{\mathbb{C}}$ would be unramified, but one knows that it is ramified at hyperelliptic curves of high genus.)

Another important difference between $\Phi_{\mathcal{N}}$ and $\Phi_{\mathcal{A}}$ is that in the case of $\Phi_{\mathcal{A}}$, by taking the union of $\Phi_{\mathcal{A}}$ and its transpose, one can compactify $\Phi_{\mathcal{A}}$ into an entirely algebraic (i.e., not just p -adic analytic)

object, namely a Hecke correspondence on \mathcal{A}_g . In the case of $\Phi_{\mathcal{N}}$, however, such a compactification into a correspondence is impossible. We refer to [Mzk4] for a detailed discussion of this phenomenon.

So far, we have been discussing the differences between $\Phi_{\mathcal{N}}$ and $\Phi_{\mathcal{A}}$. In fact, however, in one very important respect, they are very similar objects. Namely, they are both *(classical) ordinary Frobenius liftings*. A (classical) ordinary Frobenius lifting is defined as follows: Let k be a perfect field of characteristic p . Let $A \stackrel{\text{def}}{=} W(k)$ (the Witt vectors over k). Let S be a formal p -adic scheme which is formally smooth over A . Let $\Phi_S : S \rightarrow S$ be a morphism whose reduction modulo p is the Frobenius morphism. Then differentiating Φ_S defines a morphism $d\Phi_S : \Phi_S^* \Omega_{S/A} \rightarrow \Omega_{S/A}$ which is zero in characteristic p . Thus, we may form a morphism

$$\Omega_{\Phi} : \Phi_S^* \Omega_{S/A} \rightarrow \Omega_{S/A}$$

by dividing $d\Phi_S$ by p . Then Φ_S is called a *(classical) ordinary Frobenius lifting* if Ω_{Φ} is an isomorphism. Just as there is a general theory of canonical coordinates associated to real analytic Kähler metrics, there is a general theory of canonical coordinates associated to ordinary Frobenius liftings. This theory is discussed in detail in §1 of Chapter III of [Mzk1]. The main result is as follows (cf. §1 of [Mzk1], Chapter III):

Theorem 0.3. (Ordinary Frobenius Liftings) *Let $\Phi_S : S \rightarrow S$ be a (classical) ordinary Frobenius lifting. Then taking the invariants of $\Omega_{S/A}$ with respect to Ω_{Φ} gives rise to an étale local system $\Omega_{\Phi}^{\text{ét}}$ on S of free \mathbf{Z}_p -modules of rank equal to $\dim_A(S)$.*

Let $z \in S(\bar{k})$ be a point valued in the algebraic closure of k . Then $\Omega_z \stackrel{\text{def}}{=} \Omega_{\Phi}^{\text{ét}}|_z$ may be thought of as a free \mathbf{Z}_p -module of rank $\dim_A(S)$; write Θ_z for the \mathbf{Z}_p -dual of Ω_z . Let S_z be the completion of S at z . Let $\widehat{\mathbf{G}}_m$ be the completion of the multiplicative group scheme \mathbf{G}_m over $W(\bar{k})$ at 1. Then there is a unique isomorphism

$$\Gamma_z : S_z \cong \widehat{\mathbf{G}}_m \otimes_{\mathbf{Z}_p}^{\text{gp}} \Theta_z$$

such that:

(i) *the derivative of Γ_z induces the natural inclusion $\Omega_z \hookrightarrow \Omega_{S/A}|_{S_z}$;*

(ii) *the action of Φ_S on S_z corresponds to multiplication by p on $\widehat{\mathbf{G}}_m \otimes_{\mathbf{Z}_p}^{\text{gp}} \Theta_z$.*

Here, by “ $\widehat{\mathbf{G}}_m \otimes_{\mathbf{Z}_p}^{\text{gp}} \Theta_z$,” we mean the tensor product in the sense of (formal) group schemes. Thus, $\widehat{\mathbf{G}}_m \otimes_{\mathbf{Z}_p}^{\text{gp}} \Theta_z$ is noncanonically isomorphic to the product of $\dim_A(S) = \text{rank}_{\mathbf{Z}_p}(\Theta_z)$ copies of $\widehat{\mathbf{G}}_m$.

Thus, we obtain canonical multiplicative parameters on \mathcal{N} and \mathcal{C}^{ord} (from $\Phi_{\mathcal{N}}$ and $\Phi_{\mathcal{C}}$, respectively). If we apply Theorem 0.3 to $\Phi_{\mathcal{A}}$ (cf. §0.7), we obtain the Serre-Tate parameters. Moreover, note that in Theorem 0.3, the identity element “1” of the formal group scheme $\widehat{\mathbf{G}}_{\mathbf{m}} \otimes_{\mathbf{Z}_p} \Omega_z$ corresponds under Γ_z to some point $\alpha_z \in S(W(\bar{k}))$ that lifts z . That is to say,

Theorem 0.3 also gives rise to a notion of canonical liftings of points in characteristic p .

In the case of $\Phi_{\mathcal{A}}$, this notion coincides with the well-known notion of the Serre-Tate canonical lifting of an ordinary abelian variety. In the case of $\Phi_{\mathcal{N}}$, the theory of canonically lifted curves is discussed in detail in Chapter IV of [Mzk1]. In the present book, however, the theory of canonical curves in the style of Chapter IV of [Mzk1] will not play a very important role.

Remark. Certain special cases of Theorem 0.3 already appear in the work of Ihara ([Ih1-4]). In fact, more generally, the work of Ihara ([Ih1-4]) on the Schwarzian equations of Shimura curves and the possibility of constructing an analogue of Serre-Tate theory for more general hyperbolic curves anticipates, at least at a philosophical level, many aspects of the theory of [Mzk1] and the present work.

Thus, in summary, although the classical ordinary theory of [Mzk1] is not compatible with Serre-Tate theory relative to the Torelli map, it is in many respects deeply structurally analogous to Serre-Tate theory. Moreover, this close structural affinity arises from the fact that in both cases,

The ordinary locus with which the theory deals is defined by the condition that some canonical Frobenius action have slope zero.

Thus, although some readers may feel unhappy about the use of the term “ordinary” to describe the theory of [Mzk1] (i.e., despite the fact that this theory is incompatible with Serre-Tate theory), we feel that this close structural affinity arising from the common condition of a slope zero Frobenius action justifies and even renders natural the use of this terminology.

Finally, just as in the complex case, where the various indigenous bundles involved gave rise to monodromy representations of the fundamental group of the hyperbolic curve involved, in the p -adic case as well, the canonical indigenous bundle of Theorem 0.2 gives rise to a canonical Galois representation, as follows. We continue with the notation of Theorem 0.2. Let $\mathcal{N}' \rightarrow \mathcal{N}$ be the morphism $\Phi_{\mathcal{N}}$, which we think of as a covering of \mathcal{N} ; let $\mathcal{C}' \stackrel{\text{def}}{=} \mathcal{C} \otimes_{\mathcal{N}} \mathcal{N}'$. Note that \mathcal{C} and \mathcal{N}

have natural log structures (obtained by pulling back the natural log structures on $\overline{\mathcal{M}}_{g,r}$ and its tautological curve, respectively). Thus, we obtain $\mathcal{C}^{\log}, \mathcal{N}^{\log}$. Let

$$\Pi_{\mathcal{N}} \stackrel{\text{def}}{=} \pi_1(\mathcal{N}^{\log} \otimes_{\mathbf{Z}_p} \mathbf{Q}_p); \quad \Pi_{\mathcal{C}} \stackrel{\text{def}}{=} \pi_1(\mathcal{C}^{\log} \otimes_{\mathbf{Z}_p} \mathbf{Q}_p)$$

Similarly, we have $\Pi_{\mathcal{N}'}; \Pi_{\mathcal{C}'}$. Then the main result is the following (Theorem 0.4 of [Mzk1], Introduction):

Theorem 0.4. (Canonical Galois Representation) *There is a natural \mathbf{Z}_p -flat, p -adically complete “ring of additive periods” $\mathcal{D}_{\mathcal{N}}^{\text{Gal}}$ on which $\Pi_{\mathcal{N}'}$ (hence also $\Pi_{\mathcal{C}'}$ via the natural projection $\Pi_{\mathcal{C}'} \rightarrow \Pi_{\mathcal{N}'}$) acts continuously, together with a twisted homomorphism*

$$\rho : \Pi_{\mathcal{C}'} \rightarrow \text{PGL}_2(\mathcal{D}_{\mathcal{N}}^{\text{Gal}})$$

where “twisted” means with respect to the action of $\Pi_{\mathcal{C}'}$ on $\mathcal{D}_{\mathcal{N}}^{\text{Gal}}$. This representation is obtained by taking Frobenius invariants of (P, ∇_P) , using a technical tool known as crystalline induction.

Thus, in summary, the theory of [Mzk1] gives one a fairly good understanding of what happens over the ordinary locus $\overline{\mathcal{N}}_{g,r}^{\text{ord}}$, complete with analogues of various objects (monodromy representations, canonical modular coordinates, etc.) that appeared in the complex case. On the other hand, it begs the following questions:

- (1) What does the nonordinary part of $\overline{\mathcal{N}}_{g,r}$ look like? What sorts of nonordinary nilcurves can occur? In particular, what does the p -curvature of such nonordinary nilcurves look like?
- (2) Does this “classical ordinary theory” admit any sort of compactification? One sees from [Mzk4] that it does not admit any sort of compactification via correspondences. Still, since the condition of being ordinary is an “open condition,” it is natural to ask what happens to this classical ordinary theory as one goes to the boundary.

The theory of the present book answers these two questions to a large extent, not by adding on a few new pieces to [Mzk1], but by starting afresh and developing from new foundations a general theory of integral Frobenius invariant indigenous bundles. The theory of the present book will be discussed in §1.

§0.10. Intrinsic Hodge Theory

Finally, before proceeding, we discuss one more general philosophical point of view from which to view the theory of the present book. This point of view is the point of view of *intrinsic Hodge theory*. To explain what we mean by this, first let us review what is generally meant by the term “Hodge theory.” Here, by “Hodge theory,” we mean a theory giving some sort of equivalence (or at least establishing an intimate relationship between) *étale topological data* (i.e., π_1 ’s, étale cohomology, etc.) and *algebraic-geometric data* (i.e., data like differentials, cohomology groups of coherent sheaves, morphisms between varieties, etc., that exists in the purely algebraic-geometric category). For instance, p -adic Hodge theory, which relates p -adic étale cohomology groups to de Rham cohomology, is clearly a prime example of such a theory.

In our case, we would like to consider the Hodge theory of a hyperbolic curve. In fact, there are many “Hodge theories” that have been considered involving hyperbolic curves. For instance, the p -adic Hodge theory of the étale cohomology of a hyperbolic curve is one example. More generally, one can consider the “nonabelian Hodge theory” of the curve. This typically means looking at spaces of representations of the geometric fundamental group of the curve into an algebraic group G (on the étale topological side) and relating them to spaces of G -bundles with connections on the curve (on the algebraic-geometric side). In our case, however, we are interested in the *intrinsic Hodge theory* of a hyperbolic curve. By this, we mean the following:

By the term intrinsic Hodge theory of a hyperbolic curve, we mean a Hodge theory concerning the curve in which the data that appears on the algebraic-geometric side is “the curve itself.” Typically, concretely speaking, what we mean by “the curve itself” is the set of points of the curve or the moduli of the curve.

Thus, the theory of [Mzk1] or the present book may be regarded as an example of an intrinsic Hodge theory of hyperbolic curves in that it gives rise to the construction of various Galois representations that govern the moduli of a hyperbolic curve.

In fact, ultimately, the motivation for developing the intrinsic Hodge theory of hyperbolic curves came from the “anabelian philosophy of Grothendieck” (see the Introduction of [Mzk3] for more details). Let us introduce some notation. Let X_K be a hyperbolic algebraic curve over a field K of characteristic zero. Let Π_X be its fundamental group. Let \bar{K} be an algebraic closure of K , and let $\Gamma_K \stackrel{\text{def}}{=} \text{Gal}(\bar{K}/K)$. Then there is a natural exact sequence

$$1 \rightarrow \Delta_X \rightarrow \Pi_X \rightarrow \Gamma_K \rightarrow 1$$

where Δ_X is the fundamental group of $X_K \otimes_K \overline{K}$. Roughly speaking, in the terminology of the present discussion,

Grothendieck's anabelian philosophy conjectures (for certain appropriate K) the existence of an intrinsic Hodge theory in which the data on the étale topological side is the isomorphism class of the above exact sequence, while the data on the algebro-geometric side is the isomorphism class of the curve.

Thus, one of the original motivations for the development of the theory of p -adic Teichmüller theory (i.e., the theory of [Mzk1] and the present book) was to use it to realize Grothendieck's anabelian philosophy. Indeed, if one could develop some sort of Hodge theory for which the data that appears on the algebro-geometric side is the moduli of the curve, while the data appearing on the étale topological side is something involving the above exact sequence, then one should be able to use such a theory to recover the isomorphism class of X_K from the above exact sequence.




Fig. 5: Recovering a Curve as a Topological Surface Equipped with a Geometry

In fact, however, it turned out that p -adic Teichmüller theory was ill-suited for the realization of Grothendieck's anabelian philosophy. Moreover, in the meantime, the sort of result (Theorem 0.5 below) that the author had hoped to prove using p -adic Teichmüller theory was proven ([Mzk3]) using entirely different techniques. Thus, in summary,

Theorem 0.5 and p -adic Teichmüller theory represent different examples of intrinsic Hodge theories of p -adic hyperbolic curves. Moreover, these two examples generalize to the p -adic case different aspects (namely, recovering the points/isomorphism class

of the curve itself versus naturally giving rise to canonical representations of the fundamental group and canonical coordinates on the moduli space) of what (in the author's opinion) is the unique example of an intrinsic Hodge theory of hyperbolic curves in the complex case, namely, the theory (reviewed in §0.1 through §0.6) surrounding the Kôbe upper half plane uniformization.

Just as we gave illustrations of the modular aspect of the intrinsic Hodge theory of hyperbolic curves at both finite and infinite primes (i.e., Figs. 3, 4), in this § (Figs. 5,6) we give illustrations (at both finite and infinite primes) of that aspect of the intrinsic Hodge theory of hyperbolic curves that corresponds to “physically recovering the curve itself.” In the complex case, this corresponds to recovering the curve by giving on a topological surface – which is an Eilenberg-MacLane space for its fundamental group – the added structure of a “geometry” (defined by the canonical hyperbolic metric of §0.2). In fact, this geometry may be regarded – by thinking in terms of rotations about a point and flows along geodesics – as a sort of action of \mathbf{C} on the topological surface. Analogously, in the p -adic case, if we let Δ'_X denote the maximal pro- p quotient of Δ_X , then one may think of Δ'_X as defining a sort of hypothetical “ p -adic topological surface,” i.e., a (pro-)Eilenberg-MacLane space for Δ'_X . Then the “holomorphic structure” is given by the action of the Galois group Γ_K on this hypothetical p -adic topological surface. Theorem 0.5 then assures us that this is indeed just enough information to “physically recover” the curve (at least at the level of recovering its dominant points). This completes the discussion of “intrinsic Hodge theory.”

Finally, for the convenience of the reader, we state the main result of [Mzk3] (which the author had hoped to prove using p -adic Teichmüller theory). First, we introduce some notation. Fix a prime number p . Let S_K be a smooth variety over a field K of characteristic zero. Let \bar{K} be an algebraic closure of K , $\Gamma_K \stackrel{\text{def}}{=} \text{Gal}(\bar{K}/K)$. Let $\Pi_S \stackrel{\text{def}}{=} \pi_1(S_K)$ be the fundamental group of S_K (for some choice of base-point). Then if $\Delta_S \stackrel{\text{def}}{=} \pi_1(S \otimes_K \bar{K})$ is the geometric fundamental group of S_K , we have an exact sequence

$$1 \rightarrow \Delta_S \rightarrow \Pi_S \rightarrow \Gamma_K \rightarrow 1$$

Let Δ'_S be the maximal pro- p quotient of Δ_S . Then the kernel of $\Delta_S \rightarrow \Delta'_S$ is normal not only in Δ_S , but also in Π_S . Thus, we let Π'_S be the quotient of Π_S by $\text{Ker}(\Delta_S \rightarrow \Delta'_S)$.

Theorem 0.5. (Dominants Points via π_1) *Let p be a prime number. Let K be a subfield of a finitely generated field extension of \mathbf{Q}_p . Let X_K be a hyperbolic curve over K . Then for any smooth variety S_K over K , the natural map*

$$X_K(S_K)^{\text{dom}} \rightarrow \text{Hom}_{\Gamma_K}^{\text{open}}(\Pi'_S, \Pi'_X)$$

is bijective. Here $X_K(S_K)^{\text{dom}}$ denotes the set of dominant K -morphisms $S_K \rightarrow X_K$, while $\text{Hom}_{\Gamma_K}^{\text{open}}(\Pi'_S, \Pi'_X)$ denotes the set of open, continuous group homomorphisms $\Pi'_S \rightarrow \Pi'_X$ over Γ_K , considered up to composition with an inner homomorphism arising from Δ'_X .




Fig. 6: Recovering a Curve p -adically as a Galois Action

§1. Overview of the Contents of the Present Book

The purpose of §0 was to set forth as much evidence as the author could assemble in support of the claim that:

The proper p -adic analogue of the theory of the Fuchsian and Bers uniformizations should be a theory of integral Frobenius invariant indigenous bundles.

Thus, the present book – which purports to lay the foundations of a “ p -adic Teichmüller theory” – is devoted precisely to the study of (integral) Frobenius invariant indigenous bundles. In this §, we would like to summarize the contents of this book, always bearing in mind the fundamental goal of understanding and cataloguing all integral Frobenius invariant indigenous bundles.

§1.1. Major Themes

We begin by briefly discussing the major themes of the present work. These themes are shown schematically in Fig. 7.

As we saw in §0.9, the notion of a *nilcurve* – i.e., a pointed stable curve equipped with a nilpotent indigenous bundle (in positive characteristic) – is fundamental to the study of Frobenius invariant indigenous bundles over the p -adics. The nilcurves form a natural algebraic stack

$$\overline{\mathcal{N}}_{g,r}$$

which, as we saw in Theorem 0.1, is finite and flat of degree p^{3g-3+r} over $(\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$. This stack formed a natural starting point for the “classical ordinary theory” reviewed in §0.9. In fact, however, the classical ordinary theory only makes use of the *ordinary locus* of $\overline{\mathcal{N}}_{g,r}$, which, in general, is far from being the whole of (or even dense in the whole of) $\overline{\mathcal{N}}_{g,r}$. Moreover, unlike the ordinary locus of $\overline{\mathcal{N}}_{g,r}$ which is far from being proper over \mathbf{F}_p , $\overline{\mathcal{N}}_{g,r}$ itself is proper over \mathbf{F}_p , a fact which imparts a certain aura of “completeness” to $\overline{\mathcal{N}}_{g,r}$ which is lacking in $\overline{\mathcal{N}}_{g,r}^{\text{ord}}$. Thus, at any rate, if one is to have a complete theory of Frobenius invariant indigenous bundles, it is natural to start by trying to understand the stack $\overline{\mathcal{N}}_{g,r}$ in its entirety.

Unfortunately, even at the time of writing of this work, the author does not by any means have a complete understanding of $\overline{\mathcal{N}}_{g,r}$. There are still many interesting unanswered questions concerning this stack (see §2 below). At the present time, however, one does have a fairly complete understanding of the qualitative and quantitative structure of $\overline{\mathcal{N}}_{g,r}$ *generically*, i.e., of an open dense portion of $\overline{\mathcal{N}}_{g,r}$ (except for monodromy questions – see §2 below). This understanding results from studying the structure of $\overline{\mathcal{N}}_{g,r}$ in a formal neighborhood of the inverse image of a point of $(\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ corresponding to a totally degenerate curve. We shall refer to a nilcurve whose underlying curve is totally degenerate as a *molecule*. Thus,

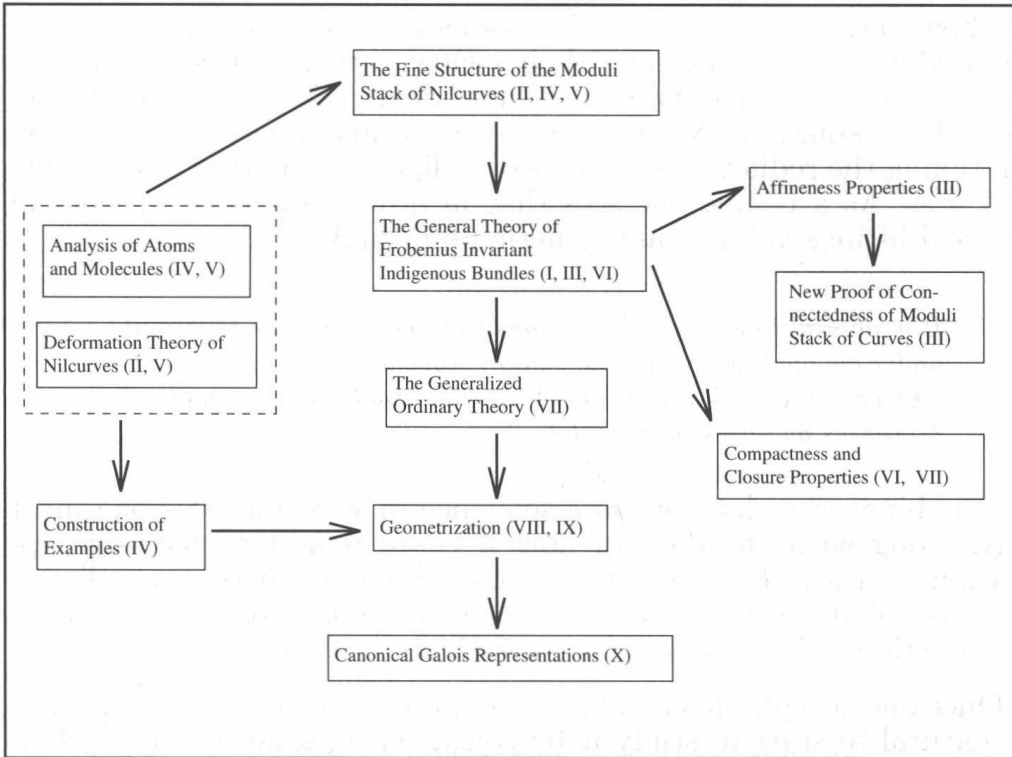


Fig. 7: The Major Themes of the Present Work.
The relevant Chapters are given in parentheses.

The first major theme of this book is the determination of the generic structure of $\overline{\mathcal{N}}_{g,r}$. This is achieved by studying the formal neighborhood in $\overline{\mathcal{N}}_{g,r}$ of all molecules, and may be broken down into two steps: (i) the complete determination of the explicit structure of all molecules; (ii) a fairly complete theory of how molecules deform.

This first major theme forms the topic of Chapters II, IV, and V.

Once one has an understanding of the generic structure of $\overline{\mathcal{N}}_{g,r}$ – which lives in characteristic p – it is natural to ask

To what extent (at least generically on $\overline{\mathcal{N}}_{g,r}$) can one obtain a theory of canonical liftings reminiscent of the classical ordinary theory?

The first obstacle that one encounters here is that if one sticks to the notion of Frobenius invariance employed in the classical ordinary theory – i.e., that the indigenous bundle be isomorphic to its own renormalized Frobenius pull-back – one has a problem: Namely, the reduction in characteristic p of any indigenous bundle satisfying this notion of invariance is *admissible*, i.e., its p -curvature has no zeroes. In

fact, however, in general, $\overline{\mathcal{N}}_{g,r}$ has generic points parametrizing nilcurves which are *dormant* (i.e., their p -curvature vanishes identically) or *spiked* (i.e., their p -curvature is not identically zero, but has zeroes, called “spikes”). Needless to say, such nilcurves could not arise from taking the reduction modulo p of indigenous bundles \mathcal{P} satisfying $\mathbf{F}^*(\mathcal{P}) \cong \mathcal{P}$. As a result, one sees that in order to obtain a theory of canonical liftings valid at any generic point of $\overline{\mathcal{N}}_{g,r}$,

It is natural to consider indigenous bundles \mathcal{P} that are invariant under the application of some (finite) sequence of \mathbf{F}^ ’s (renormalized Frobenius pull-backs) and Φ^* ’s (the usual non-renormalized Frobenius pull-back of crystals).*

The combinatorial data of *which* sequence of \mathbf{F}^* ’s and Φ^* ’s one must apply, along with the data of how \mathcal{P} changes under successive application of these \mathbf{F}^* ’s and Φ^* ’s is contained in an object called a *VF-pattern*. Thus, *each VF-pattern determines a specific notion of Frobenius invariance.* (Here, “VF” stands for “Verschiebung-Frobenius.”)

Once one accepts this more general notion of Frobenius invariance, it is natural to start to study it by constructing some sort of *algebraic moduli stack of indigenous bundles invariant in the fashion prescribed by some fixed VF-pattern*. Thus, for each VF-pattern Π , we obtain an algebraic stack

$$\mathcal{Q}^\Pi$$

Once one has this stack, it is natural to try to understand its global structure:

The main results concerning \mathcal{Q}^Π are: (i) results to the effect that it tends to be affine or, at least, quasi-affine; (ii) results to the effect that it tends to be closed and disjoint from the closure of other $\mathcal{Q}^{\Pi'}$ (for VF-patterns $\Pi' \neq \Pi$) inside the space of all indigenous bundles. Moreover, when interpreted properly, these affineness results give a new proof of the connectedness of $\mathcal{M}_{g,r}$.

The results just mentioned form the topic of Chapters I, III, and VI. Of particular interest relative to the asserted analogy between complex Teichmüller theory and the theory of the present book is the fact that

Just as complex Teichmüller theory shows that $\mathcal{M}_{g,r}$ is connected by showing that Teichmüller space is contractible, the theory of the present book allows one to derive the connectedness of $\mathcal{M}_{g,r}$ from the “crystalline contractibility” of (certain) \mathcal{Q}^Π .

Although, historically, many different proofs have been given of the connectedness of $\mathcal{M}_{g,r}$, the author believes that this proof is the first essentially characteristic p proof of this result.

Unfortunately, if one studies the stack \mathcal{Q}^Π in this generality, it is difficult to say much more concerning its explicit structure. If, however, one imposes an appropriate notion of *ordinariness* – thus giving rise to an open substack $(\mathcal{Q}^\Pi)^{\text{ord}}$ of \mathcal{Q}^Π – which is analogous to the notion discussed in [Mzk1], but depends on the VF-pattern Π , then one can say much more. This is the topic of Chapter VII. Namely,

Not only does one have a much more explicit understanding of the structure of $(\mathcal{Q}^\Pi)^{\text{ord}}$ than of \mathcal{Q}^Π , one also obtains in the “generalized ordinary case” (just as in the theory of [Mzk1]) a system of canonical Frobenius liftings on a certain canonical p-adic formal stack which is étale over $(\overline{\mathcal{M}}_{g,r})_{\mathbf{Z}_p}$. Moreover, for certain types of VF-pattern, one can prove that $(\mathcal{Q}^\Pi)^{\text{ord}}$ is not only open, but also “ ω -closed” (roughly speaking, “closed as far as the differentials are concerned” – cf. Chapter VII for more details) in \mathcal{Q}^Π .

Of course in the theory of [Mzk1], one gets a single canonical Frobenius lifting rather than a whole system of canonical Frobenius liftings, as we do in the present “generalized ordinary context.” It turns out, however, that in the present generalized ordinary context, such a *system* (as opposed to a single lifting) occurs very naturally and is the proper analogue of the single lifting of the theory of [Mzk1]. Next, let us remark that the fact that (for certain types of Π) $(\mathcal{Q}^\Pi)^{\text{ord}}$ is not only open, but also ω -closed in \mathcal{Q}^Π is of particular interest in that:

One of the main motivations for studying “the rest of” (i.e., the nonordinary portion of) $\overline{\mathcal{N}}_{g,r}$ and, indeed, for the entire theory of the present book was to “compactify” the theory of [Mzk1]. That is to say, because the theory of [Mzk1] is only defined over an open substack of $\overline{\mathcal{N}}_{g,r}$, one has a sense that it is, in some sense, not “complete.” The interesting and perhaps surprising conclusion that was eventually reached, however, is that inside \mathcal{Q}^Π , or indeed, inside the space of all indigenous bundles, the theory of [Mzk1] is, in some sense, already complete!

Looking back at the theory of [Mzk1] as reviewed in §0.9, so far we have obtained a generalized analogue of Theorem 0.2. Thus, the next natural step is to ask to what extent Theorem 0.3 – which allowed us to derive from the canonical Frobenius lifting a canonical local uniformization by the multiplicative group $\widehat{\mathbf{G}}_m$ – can be extended to the generalized ordinary context. Namely, one wants to know to what extent the canonical *system* of Frobenius liftings discussed above gives rise to some sort of canonical local uniformization. Unfortunately, we

are not able to achieve this for an arbitrary VF-pattern. In the following, we shall refer to the process of converting a system of Frobenius liftings into a canonical local uniformization as the *geometrization* of the system of Frobenius liftings. What we are able to achieve is the following:

For certain fairly general types of VF-pattern, the resulting system of Frobenius liftings can be geometrized. However, unlike the classical ordinary case, instead of obtaining uniformizations by $\widehat{\mathbf{G}}_m$, one obtains uniformizations by more general types of Lubin-Tate groups, twisted products of Lubin-Tate groups, and fibrations of certain such twisted products by other such twisted products.

These exotic new p -adic geometries are discussed in more detail below (§1.6).

Incidentally, we note one more important theme of this book which is relevant here: Namely, these geometries occur on certain p -adic formal stacks that are étale over $(\overline{\mathcal{M}}_{g,r})_{\mathbf{Z}_p}$. Thus, in order to know that the theory of such geometries is nonvacuous, it is important to know that these formal stacks are *nonempty*. It turns out that, unlike in the classical ordinary case, where it is fairly trivial to show that $\overline{\mathcal{N}}_{g,r}^{\text{ord}}$ is nonempty, in the generalized ordinary context, it is highly nontrivial in general to show that the resulting stack is nonempty. However, by combining the deformation theory of molecules with certain other technical tools, we are able to construct a number of explicit examples (Chapter IV) which show that in a fairly large and representative number of cases, the relevant stack is nonempty.

Finally, recall that the last major portion of the theory of [Mzk1] (as reviewed in §0.9) is the theory of the *canonical Galois representation*, i.e., Theorem 0.4. It turns out that in the cases where geometrization is possible, this theory generalizes in a fairly straightforward fashion to the generalized ordinary context. (Just as in the classical ordinary case, the key technique here is the systematic application of the technical tool called *crystalline induction* in [Mzk1].) This portion of the theory is philosophically important in that

Any p -adic theory that purports to be the p -adic analogue of the theory of the Fuchsian uniformization in the complex case should ultimately give rise to representations which are analogous to the canonical representation of the topological fundamental group into $\text{PSL}_2(\mathbf{R})$ that one obtains in the complex case.

Chapter X assures us that this is indeed the case.

In the remainder of this §, we discuss these major themes in more detail, giving, where possible, explicit statements of theorems that realize each of these major themes.

§1.2. Atoms, Molecules, and Nilcurves

Let p be an odd prime. Let g and r be nonnegative integers such that $2g - 2 + r \geq 1$. Let $\overline{\mathcal{N}}_{g,r}$ be the stack of nilcurves in characteristic p . We denote by $\mathcal{N}_{g,r} \subseteq \overline{\mathcal{N}}_{g,r}$ the open substack consisting of *smooth nilcurves*, i.e., nilcurves whose underlying curve is smooth. Then the first step in our analysis of $\overline{\mathcal{N}}_{g,r}$ is the introduction of the following notions (cf. Definitions 1.1 and 3.1 of Chapter II in the text):

Definition 1.1. We shall call a nilcurve *dormant* if its p -curvature (i.e., the p -curvature of its underlying indigenous bundle) is identically zero. Let d be a nonnegative integer. Then we shall call a smooth nilcurve *spiked of strength d* if the zero locus of its p -curvature forms a divisor of degree d .

If d is a nonnegative integer (respectively, the symbol ∞), then we shall denote by

$$\mathcal{N}_{g,r}[d] \subseteq \mathcal{N}_{g,r}$$

the locally closed substack of nilcurves that are spiked of strength d (respectively, dormant). It is immediate that there does indeed exist such a locally closed substack, and that if k is an algebraically closed field of characteristic p , then

$$\mathcal{N}_{g,r}(k) = \coprod_{d=0}^{\infty} \mathcal{N}_{g,r}[d](k)$$

Moreover, we have the following result (cf. Chapter II, Theorems 1.12, 2.8, and 3.9):

Theorem 1.2. (Stratification of $\mathcal{N}_{g,r}$) *Any two irreducible components of $\overline{\mathcal{N}}_{g,r}$ intersect. Moreover, for $d = 0, 1, \dots, \infty$, the stack $\mathcal{N}_{g,r}[d]$ is smooth over \mathbf{F}_p of dimension $3g - 3 + r$ (if it is nonempty). Finally, $\mathcal{N}_{g,r}[\infty]$ is irreducible, and its closure in $\overline{\mathcal{N}}_{g,r}$ is smooth over \mathbf{F}_p .*

Thus, in summary, we see that

The classification of nilcurves by the size of the zero locus of their p -curvatures induces a natural decomposition of $\mathcal{N}_{g,r}$ into smooth (locally closed) strata.

Unfortunately, however, Theorem 1.2 still only gives us a very rough idea of the structure of $\mathcal{N}_{g,r}$. For instance, it tells us nothing of the degree of each $\mathcal{N}_{g,r}[d]$ over $\mathcal{M}_{g,r}$.

Remark. Some people may object to the use of the term “stratification” here for the reason that in certain contexts (e.g., the Ekedahl-Oort stratification of the moduli stack of principally polarized abelian varieties – cf. [Geer], §2), this term is only used for decompositions into locally closed subschemes whose closures satisfy certain (rather stringent) axioms. Here, we do not mean to imply that we can prove any nontrivial results concerning the closures of the $\mathcal{N}_{g,r}[d]$ ’s. That is to say, in this book, we will use the term “stratification” only in the weak sense (i.e., that $\mathcal{N}_{g,r}$ is the union of the $\mathcal{N}_{g,r}[d]$). This usage conforms to the usage of Lecture 8 of [Surf], where “flattening stratifications” are discussed.

In order to understand things more explicitly, it is natural to attempt to do the following:

- (1) *Understand the structure – especially, what the p -curvature looks like – of all molecules (i.e., nilcurves whose underlying curve is totally degenerate).*
- (2) *Understand how each molecule deforms, i.e., given a molecule, one can consider its formal neighborhood \mathcal{N} in $\overline{\mathcal{N}}_{g,r}$. Then one wants to know the degree of each $\mathcal{N} \cap \mathcal{N}_{g,r}[d]$ (for all d) over the corresponding formal neighborhood \mathcal{M} in $\overline{\mathcal{M}}_{g,r}$.*

Obtaining a complete answer to these two questions is the topic of Chapters IV and V.

First, we consider the problem of understanding the structure of *molecules*. Since the underlying curve of a molecule is a totally degenerate curve – i.e., a stable curve obtained by gluing together \mathbf{P}^1 ’s with three nodal/marked points – it is natural to restrict the given nilpotent indigenous bundle on the whole curve to each of these \mathbf{P}^1 ’s with three marked points. Thus, for each irreducible component of the original curve, we obtain a \mathbf{P}^1 with three marked points equipped with something very close to a nilpotent indigenous bundle. The only difference between this bundle and an indigenous bundle is that its monodromy at some of the marked points (i.e., those marked points that correspond to nodes on the original curve) might not be nilpotent. In general, a bundle (with connection) satisfying all the conditions that an indigenous bundle satisfies except that its monodromy at the marked points might not be nilpotent is called a *torally indigenous bundle* (cf. Chapter I, Definition 4.1). (When there is fear of confusion, indigenous bundles in the strict sense (as in [Mzk1], Chapter I) will be called *classical indigenous*.) For simplicity, we shall refer to any pointed stable curve (respectively, totally degenerate pointed stable curve) equipped with a nilpotent *torally indigenous* bundle as a *nilcurve* (respectively, *molecule*) (cf. §0 of Chapter V). Thus, when it is necessary to avoid confusion with the toral case, we shall say that “ $\overline{\mathcal{N}}_{g,r}$ is

the stack of classical nilcurves.” Finally, we shall refer to a (possibly toral) nilcurve whose underlying curve is \mathbf{P}^1 with three marked points as an *atom*.

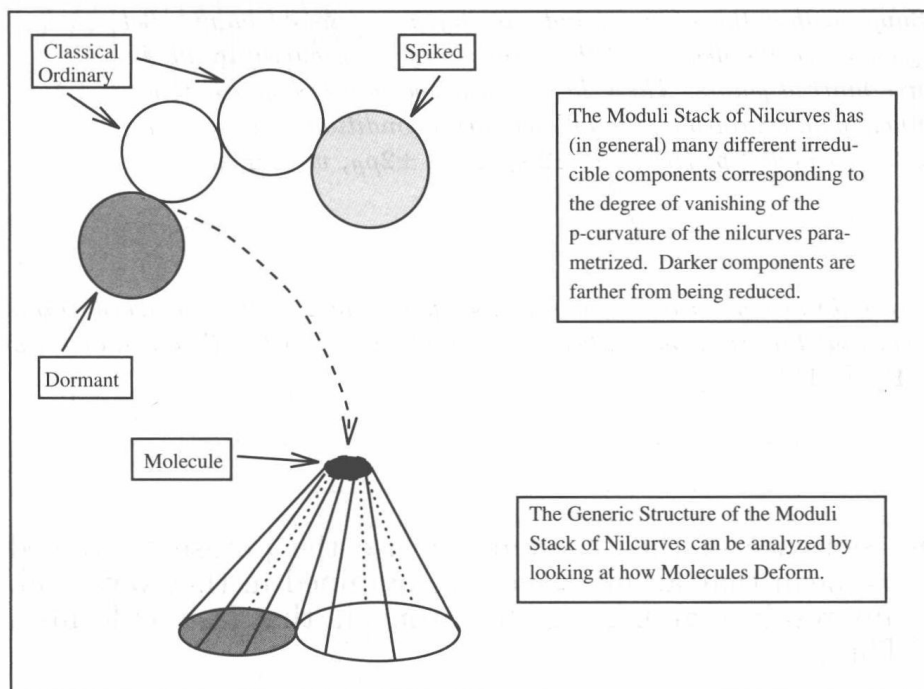


Fig. 8: The Structure of $\overline{\mathcal{N}}_{g,r}$

At any rate, to summarize, a molecule may be regarded as being made of atoms. It turns out that the monodromy at each marked point of an atom (or, in fact, more generally any nilcurve) has an invariant called the *radius*. The radius is, strictly speaking, an element of $\mathbf{F}_p/\{\pm 1\}$ (cf. Proposition 1.5 of Chapter II) – i.e., the quotient set of \mathbf{F}_p by the action of ± 1 – but, by abuse of notation, we shall often speak of the radius ρ as an element of \mathbf{F}_p . Then we have the following answer to (1) above (cf. §1 of Chapter V):

Theorem 1.3. (The Structure of Atoms and Molecules) *The structure theory of atoms (over any field of characteristic p) may be summarized as follows:*

- (1) *The three radii of an atom define a bijection of the set of isomorphism classes of atoms with the set of ordered triples of elements of $\mathbf{F}_p/\{\pm 1\}$.*

(2) For any triple of elements $\rho_\alpha, \rho_\beta, \rho_\gamma \in \mathbf{F}_p$, there exist integers $a, b, c \in [0, p-1]$ such that (i) $a \equiv \pm 2\rho_\alpha, b \equiv \pm 2\rho_\beta, c \equiv \pm 2\rho_\gamma$; (ii) $a+b+c$ is odd and $< 2p$. Moreover, the atom of radii $\rho_\alpha, \rho_\beta, \rho_\gamma$ is dormant if and only if the following three inequalities are satisfied simultaneously: $a+b > c, a+c > b, b+c > a$.

(3) Suppose that the atom of radii $\rho_\alpha, \rho_\beta, \rho_\gamma$ is nondormant. Let $v_\alpha, v_\beta, v_\gamma$ be the degrees of the zero loci of the p -curvature at the three marked points. Then the nonnegative integers $v_\alpha, v_\beta, v_\gamma$ are uniquely determined by the following two conditions: (i) $v_\alpha + v_\beta + v_\gamma$ is odd and $< p$; (ii) $v_\alpha \equiv \pm 2\rho_\alpha, v_\beta \equiv \pm 2\rho_\beta, v_\gamma \equiv \pm 2\rho_\gamma$.

Molecules are obtained precisely by gluing together atoms at their marked points in such a way that the radii at marked points that are glued together coincide (as elements of $\mathbf{F}_p/\{\pm 1\}$).

In the last sentence of the theorem, we use the phrase “obtained precisely” to mean that all molecules are obtained in that way, and, moreover, any result of gluing together atoms in that fashion forms a molecule. Thus,

Theorem 1.3 reduces the structure theory of atoms and molecules to a matter of combinatorics.

Theorem 1.3 follows from the theory of Chapter IV.

Before proceeding, we would like to note the analogy with the theory of “pants” (see [Abik] for an exposition) in the complex case. In the complex case, the term “pants” is used to describe a Riemann surface which is topologically isomorphic to a Riemann sphere minus three points. The *holomorphic* isomorphism class of such a Riemann surface is given precisely by specifying three radii, i.e., the size of its three holes. Moreover, any hyperbolic Riemann surface can be analyzed by decomposing it into a union of pants, glued together at the boundaries. Thus, there is a certain analogy between the theory of pants and the structure theory of atoms and molecules given in Theorem 1.3.

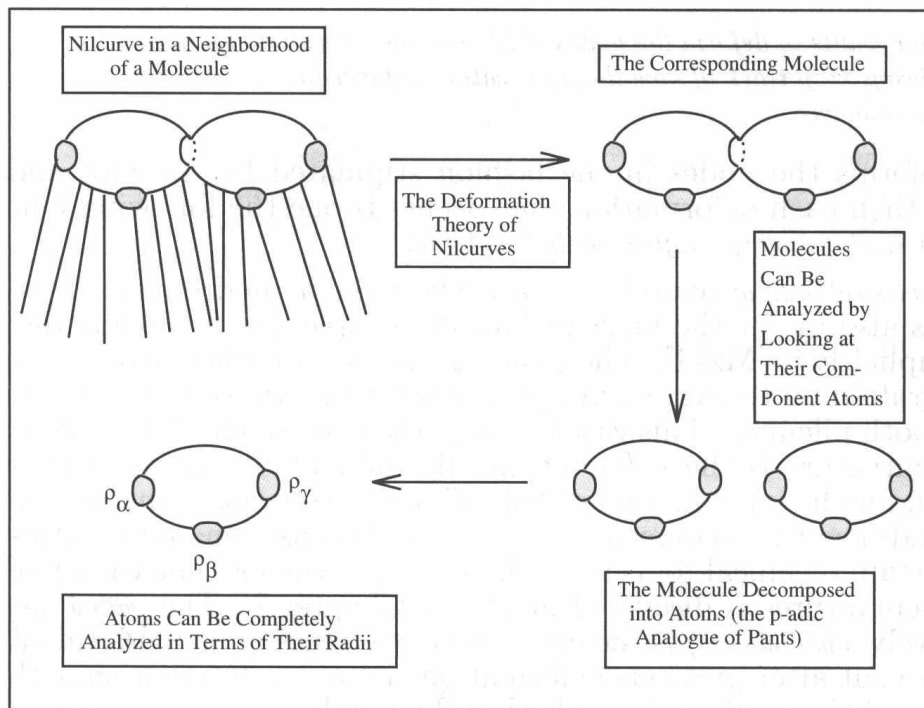


Fig. 9: The Steps Used to Analyze the Structure of $\overline{\mathcal{N}}_{g,r}$

Next, we turn to the issue of understanding how molecules deform. Let M be a nondormant *classical molecule* (i.e., it has nilpotent monodromy at all the marked points). Let us write n_{tor} for the number of “toral nodes” (i.e., nodes at which the monodromy is not nilpotent) of M . Let us write n_{dor} for the number of dormant atoms in M . To describe the deformation theory of M , it is useful to choose a *plot* Π for M . A plot is an ordering satisfying certain conditions) of a certain subset of the nodes of M (see §1 of Chapter V for more details). This ordering describes the order in which one deforms the nodes of M . (Despite the similarity in notation, plots have nothing to do with VF-patterns.) Once the plot is fixed, one can contemplate the various *scenarios* that may occur. Roughly speaking, a scenario is an assignment (satisfying certain conditions) of one of the three symbols $\{0, +, -\}$ to each of the branches of each of the nodes of M (see §1 of Chapter V for more details). There are $2^{n_{\text{dor}}}$ possible scenarios.

The point of all this terminology is the following:

One wants to deform the nodes of M in a such a way that one can always keep track of how the p -curvature deforms as one deforms the nilcurve.

If one deforms the nodes in the fashion stipulated by the plot and scenario, then each deformation that occurs is one the following four types: *classical ordinary*, *grafted*, *philial*, *aphilial*.

The *classical ordinary* case is the case where the monodromy is nilpotent. It is also by far the most technically simple and is already discussed implicitly in [Mzk1]. The *grafted case* is the case where a dormant atom is grafted on to (what after previous deformations is) a nondormant smooth nilcurve. This case is where the consequent deformation of the p -curvature is the most technically difficult to analyze and is the reason for the introduction of “plots” and “scenarios.” In order to understand how the p -curvature deforms in this case, one must introduce a certain technical tool called the *virtual p -curvature*. The theory of virtual p -curvatures is discussed in §2.2 of Chapter V. The *philial case* (respectively, *aphilial case*) is the case where one glues on a nondormant atom to (what after previous deformations is) a nondormant smooth nilcurve, and the parities (i.e., whether the number is even or odd) of the vanishing orders of the p -curvature at the two branches of the node are *opposite* to one another (respectively, the *same*). In the philial case (respectively, aphilial case), deformation gives rise to a spike (respectively, no spike). An illustration of these four fundamental types of deformation is given in Fig. 10. The signs in this illustration are the signs that are assigned to the branches of the nodes by the “scenario.” When the p -curvature is not identically zero (i.e., on the light-colored areas), this sign is the parity (i.e., plus for even, minus for odd) of the vanishing order of the p -curvature. For a given scenario Σ , we denote by $n_{\text{phl}}(\Sigma)$ (respectively, $n_{\text{aph}}(\Sigma)$) the number of philial (respectively, aphilial) nodes that occur when the molecule is deformed according to that scenario.

If $U = \text{Spec}(A)$ is a connected noetherian scheme of dimension 0, then we shall refer to the length of the artinian ring A as the *padding degree* of U . Then the theory just discussed gives rise to the following answer to (2) above (cf. Theorem 1.1 of Chapter V):

Theorem 1.4. (Deformation Theory of Molecules) *Let M be a classical molecule over an algebraically closed field k of characteristic p . Let \mathcal{N} be the completion of $\overline{\mathcal{N}}_{g,r}$ at M . Let \mathcal{M} be the completion of $(\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ at the point defined by the curve underlying M . Let $\overline{\eta}$ be the strict henselization of the generic point of \mathcal{M} . Then the natural morphism $\mathcal{N} \rightarrow \mathcal{M}$ is finite and flat of degree $2^{n_{\text{tor}}}$. Moreover:*

- (1) *If M is dormant, then $\mathcal{N}_{\text{red}} \cong \mathcal{M}$, and $\mathcal{N}_{\overline{\eta}}$ has padding degree 2^{3g-3+r} .*

- (2) If M is nondormant, fix a plot Π for M . Then for each of the $2^{n_{\text{dor}}}$ scenarios associated to Π , there exists a natural open substack $\mathcal{N}_\Sigma \subseteq \mathcal{N}_{\bar{\eta}} \stackrel{\text{def}}{=} \mathcal{N} \times_{\mathcal{M}} \bar{\eta}$ such that: (i.) $\mathcal{N}_{\bar{\eta}}$ is the disjoint union of the \mathcal{N}_Σ (as Σ ranges over all the scenarios); (ii.) every residue field of \mathcal{N}_Σ is separable over (hence equal to) $k(\bar{\eta})$; (iii) the degree of $(\mathcal{N}_\Sigma)_{\text{red}}$ over $\bar{\eta}$ is $2^{n_{\text{aph}}(\Sigma)}$; (iv) each connected component of \mathcal{N}_Σ has padding degree $2^{n_{\text{phl}}(\Sigma)}$; (v) the smooth nil-curve represented by any point of $(\mathcal{N}_\Sigma)_{\text{red}}$ is spiked of strength $p \cdot n_{\text{phl}}(\Sigma)$.

In particular, this Theorem reduces the computation of the degree of any $\mathcal{N}_{g,r}[d]$ over $(\mathcal{M}_{g,r})_{\mathbb{F}_p}$ to a matter of combinatorics.




Fig. 10: The Four Types of Nodal Deformation

For instance, let us denote by $n_{g,r,p}^{\text{ord}}$ the degree of $\mathcal{N}_{g,r}^{\text{ord}}$ (which – as a consequence of Theorem 1.4! (cf. Corollary 1.2 of Chapter V) – is open and dense in $\mathcal{N}_{g,r}[0]$) over $(\mathcal{M}_{g,r})_{\mathbb{F}_p}$. Then following the algorithm implicit in Theorem 1.4, $n_{g,r,p}^{\text{ord}}$ is computed explicitly for low g and r in Corollary 1.3 of Chapter V. Moreover, we note the following two interesting phenomena:

- (1) Degrees such as $n_{g,r,p}^{\text{ord}}$ tend to be well-behaved – even polynomial, with coefficients equal to various integrals over Euclidean space – as functions of p . Thus, for instance, the limit, as p goes to infinity, of $n_{0,r,p}^{\text{ord}}/p^{r-3}$ exists and is equal to the volume of a certain polyhedron in $(r-3)$ -dimensional Euclidean space. See Corollary 1.3 of Chapter V for more details.
- (2) Theorem 1.4 gives, for every choice of totally degenerate r -pointed stable curve of genus g , an (a priori) distinct algorithm for computing $n_{g,r,p}^{\text{ord}}$. Since $n_{g,r,p}^{\text{ord}}$, of course, does not depend on the choice of underlying totally degenerate curve, we thus obtain equalities between the various sums that occur (to compute $n_{g,r,p}^{\text{ord}}$) in each case. If one writes out these equalities, one thus obtains various combinatorial identities. Although the author has yet to achieve a systematic understanding of these combinatorial identities, already in the cases that have been computed (for low g and r), these identities reduce to such nontrivial combinatorial facts as Lemmas 3.5 and 3.6 of Chapter V.

Although the author does not have even a conjectural theoretical understanding of these two phenomena, he nonetheless feels that they are very interesting and deserve further study.

§1.3. The \mathcal{MF}^∇ -Object Point of View

Before discussing the general theory of canonical liftings of nilpotent indigenous bundles, it is worth stopping to examine the general conceptual context in which this theory will be developed. To do this, let us first recall the theory of \mathcal{MF}^∇ -objects developed in §2 of [Falt1]. Let p be a prime number, and let S be a smooth \mathbf{Z}_p -scheme. Then in *loc. cit.*, a certain category $\mathcal{MF}^\nabla(S)$ is defined. Objects of this category $\mathcal{MF}^\nabla(S)$ consist of

- (1) a vector bundle \mathcal{E} on S equipped with an integrable connection $\nabla_{\mathcal{E}}$; one may equivalently regard the pair $(\mathcal{E}, \nabla_{\mathcal{E}})$ as a crystal on the crystalline site $\text{Crys}(S \otimes_{\mathbf{Z}_p} \mathbf{F}_p/\mathbf{Z}_p)$ valued in the category of vector bundles;
- (2) a filtration $F^\bullet(\mathcal{E}) \subseteq \mathcal{E}$ (called the Hodge filtration) of subbundles of \mathcal{E} ;
- (3) a Frobenius action $\Phi_{\mathcal{E}}$ on the crystal $(\mathcal{E}, \nabla_{\mathcal{E}})$;

moreover, these objects satisfy certain conditions, which we omit here. (Strictly speaking, here we are only considering those objects of $\mathcal{MF}^\nabla(S)$ as defined in *loc. cit.* which are flat over \mathbf{Z}_p .) Let Π_S be the fundamental group of $S \otimes_{\mathbf{Z}_p} \mathbf{Q}_p$ (for some choice of base-point). In *loc. cit.*, for each $\mathcal{MF}^\nabla(S)$ -object $(\mathcal{E}, \nabla_{\mathcal{E}}, F^\bullet(\mathcal{E}), \Phi_{\mathcal{E}})$, a certain natural Π_S -module V is constructed by taking invariants of $(\mathcal{E}, \nabla_{\mathcal{E}})$ with respect to its Frobenius action $\Phi_{\mathcal{E}}$. If \mathcal{E} is of rank r , then V is a free \mathbf{Z}_p -module of rank r . The most fundamental example of an $\mathcal{MF}^\nabla(S)$ -object is the following:

If $X \rightarrow S$ is a smooth proper morphism, then the relative de Rham cohomology of $X \rightarrow S$ forms an $\mathcal{MF}^\nabla(S)$ -object. It is shown, moreover, in [Falt1], that the Π_S -module defined by this $\mathcal{MF}^\nabla(S)$ -object is the Π_S -module given by the corresponding relative étale cohomology module of $X \rightarrow S$.

For instance, if S is the moduli stack of principally polarized abelian varieties, and $X \rightarrow S$ is the tautological abelian variety, then the relative first de Rham cohomology module of $X \rightarrow S$ forms an $\mathcal{MF}^\nabla(S)$ -module whose restriction to the ordinary locus of S is (by Serre-Tate theory) intimately connected to the “ p -adic uniformization theory” of S .

In the context of the present work, we would like to consider the case where $S = (\mathcal{M}_{g,r})_{\mathbf{Z}_p}$. Moreover, just as the first de Rham cohomology module of the universal abelian variety gives rise to a “fundamental uniformizing $\mathcal{MF}^\nabla(S)$ -module” on the moduli stack of principally polarized abelian varieties, we would like to define and study a corresponding “fundamental uniformizing \mathcal{MF}^∇ -object” on $(\mathcal{M}_{g,r})_{\mathbf{Z}_p}$. Unfortunately, as long as one sticks to the conventional definition of \mathcal{MF}^∇ -object given in [Falt1], it appears that such a natural “fundamental uniformizing \mathcal{MF}^∇ -object” simply does not exist on $(\mathcal{M}_{g,r})_{\mathbf{Z}_p}$. This is not so surprising in view of the nonlinear nature of the Teichmüller group (i.e., the fundamental group of $(\mathcal{M}_{g,r})_{\mathbf{C}}$). In order to obtain a natural “fundamental uniformizing \mathcal{MF}^∇ -object” on $(\mathcal{M}_{g,r})_{\mathbf{Z}_p}$, one must generalize the “classical” linear notion of [Falt1] as follows: Instead of considering crystals (equipped with filtrations and Frobenius actions) valued in the category of *vector bundles*, one must consider crystals (still equipped with filtrations and Frobenius actions in some appropriate sense) valued in the category of *schemes* (or more generally, *algebraic spaces*). Thus,

One philosophical point of view from which to view the present work is that it is devoted to the study of a certain canonical uniformizing \mathcal{MF}^∇ -object on $(\mathcal{M}_{g,r})_{\mathbf{Z}_p}$ valued in the category of algebraic spaces.

Just as in the case of abelian varieties, this canonical uniformizing

\mathcal{MF}^∇ -object will be obtained by taking some sort of de Rham cohomology of the universal curve over $(\mathcal{M}_{g,r})_{\mathbf{Z}_p}$. The rest of this subsection is devoted to describing this \mathcal{MF}^∇ -object in more detail.

Now let S be the spectrum of an algebraically closed field (of characteristic not equal to 2), and let X be a smooth, proper, geometrically curve over S of genus ≥ 2 . Let $P \rightarrow X$ be a \mathbf{P}^1 -bundle equipped with a connection ∇_P . If $\sigma : X \rightarrow P$ is a section of this \mathbf{P}^1 -bundle, then we shall refer to the number $\frac{1}{2}\deg(\sigma^*\tau_{P/X})$ (where $\tau_{P/X}$ is the relative tangent bundle of P over X) as the *canonical height* of σ . Moreover, note that by differentiating σ by means of ∇_P , one obtains a morphism $\tau_X \rightarrow \sigma^*\tau_{P/X}$. We shall say that σ is *horizontal* if this morphism is identically zero. Then we make the following definitions:

- (i) (*Roughly speaking*) we shall call (P, ∇_P) *crys-stable* if it does not admit any horizontal sections of canonical height ≤ 0 . (See Definition 1.2 of Chapter I for a precise definition.)
- (ii) (*Roughly speaking*) we shall call (P, ∇_P) *crys-stable of level 0* (or just *stable*) if it does not admit any sections of canonical height ≤ 0 . (See Definition 3.2 of Chapter I for a precise definition.)
- (iii) Let l be a positive half-integer (i.e., a positive element of $\frac{1}{2}\mathbf{Z}$). We shall call (P, ∇_P) *crys-stable of level l* if it admits a section of canonical height $-l$. If it does admit such a section, then this section is the unique section of $P \rightarrow X$ of negative canonical height. This section will be referred to as the *Hodge section*. (See Definition 3.2 of Chapter I for more details.)

For instance, if \mathcal{E} is a vector bundle of rank two on X such that $\text{Ad}(\mathcal{E})$ is a stable vector bundle on X (of rank three), and $P \rightarrow X$ is the projective bundle associated to \mathcal{E} , then (P, ∇_P) will be *crys-stable of level 0* (regardless of the choice of ∇_P). On the other hand, an indigenous bundle on X will be *crys-stable of level $g - 1$* . More generally, these definitions generalize to the case when X is a family of pointed stable curves over an arbitrary base (on which 2 is invertible).

Now we are ready to begin to describe the nonlinear \mathcal{MF}^∇ -object on $(\mathcal{M}_{g,r})_{\mathbf{Z}_p}$ (where p is odd) that will be the topic of the present work:

- (1) **The crystal in algebraic spaces:** There is a natural *fine moduli space* $\mathcal{Y} \rightarrow (\mathcal{M}_{g,r})_{\mathbf{Z}_p}$ of *crys-stable bundles* on the universal curve. Here, \mathcal{Y} is a relative algebraic space over $(\mathcal{M}_{g,r})_{\mathbf{Z}_p}$ of dimension $2(3g - 3 + r)$. Moreover, \mathcal{Y} is equipped with a connection $\nabla_{\mathcal{Y}}$ (in the Grothendieck sense: i.e., an isomorphism of the two pull-backs of \mathcal{Y}

to the first infinitesimal neighborhood of the diagonal in $(\mathcal{M}_{g,r})_{\mathbf{Z}_p} \times_{\mathbf{Z}_p} (\mathcal{M}_{g,r})_{\mathbf{Z}_p}$ which is the identity on the diagonal). Thus, $(\mathcal{Y}, \nabla_{\mathcal{Y}})$ defines a crystal in algebraic spaces on $(\mathcal{M}_{g,r})_{\mathbf{Z}_p}$. (See Theorem 2.7, Proposition 3.1 of Chapter I for more details.)

- (2) **The Hodge structure:** In the present nonlinear context, the analogue of the Hodge filtration is a collection of subspaces of the space \mathcal{Y} . The subspaces considered are the *fine moduli spaces* \mathcal{Y}^l of *crys-stable bundles of level l* . Here, \mathcal{Y}^0 is open in \mathcal{Y} , while the subspace \mathcal{Y}^χ (where $\chi = \frac{1}{2}(2g - 2 + r)$) – consisting of the indigenous bundles on the universal curve – is closed in \mathcal{Y} . (See Chapter I, Proposition 3.3, Lemmas 3.4 and 3.8, and Theorem 3.10 for more details.)

Remark. This collection of subspaces is reminiscent of the stratification (on the moduli stack of smooth nilcurves) of §1.2. This is by no means a mere coincidence. In fact, in some sense, the stratification of $\mathcal{N}_{g,r}$ which was discussed in §1.2 is the Frobenius conjugate of the Hodge structure mentioned above. That is to say, the relationship between these two collections of subspaces is the nonlinear analogue of the relationship between the filtration on the de Rham cohomology of a variety in positive characteristic induced by the “conjugate spectral sequence” and the Hodge filtration on the cohomology. (That is to say, the former filtration is the Frobenius conjugate of the latter filtration.)

At any rate, to summarize:

The algebraic portion (i.e., crystal in algebraic spaces plus Hodge structure) of the canonical \mathcal{MF}^∇ -object that we will consider on $(\mathcal{M}_{g,r})_{\mathbf{Z}_p}$ is given by considering spaces of crys-stable bundles. Put another way, it is a sort of de Rham-theoretic H^1 with coefficients in PGL_2 of the universal curve over $\mathcal{M}_{g,r}$.

The purely algebraic theory of crys-stable bundles is the topic of Chapter I of the present work. This purely algebraic theory is by far the easiest aspect of the canonical \mathcal{MF}^∇ -object under consideration. In some sense, the rest of the book (i.e., the complement of Chapter I) is devoted to the much less trivial task of defining and studying the Frobenius action. The issue of defining this Frobenius action will be discussed in the following subsection.

§1.4. The Generalized Notion of a Frobenius Invariant Indigenous Bundle

In this subsection, we would like to take up the task of describing the Frobenius action on crys-stable bundles. Just as in the case of the linear \mathcal{MF}^∇ -objects of [Falt1], and as motivated by comparison with the complex case (see the discussion of §0), we are interested in *Frobenius invariant sections of the \mathcal{MF}^∇ -object*, i.e., Frobenius invariant bundles. Moreover, since ultimately we are interested in *uniformization theory*, instead of studying general Frobenius invariant crys-stable bundles, we will only consider *Frobenius invariant indigenous bundles*. The reason that we must nonetheless introduce crys-stable bundles is that (cf. the discussion of §1.1) in order to obtain canonical lifting theories that are valid at generic points of $\mathcal{N}_{g,r}$ parametrizing *dormant* or *spiked* nilcurves, it is necessary to consider indigenous bundles that are fixed not (necessarily) after *one* application of Frobenius, but after *several* applications of Frobenius. As one applies Frobenius over and over again, the bundles that appear at intermediate stages need not be indigenous. They will, however, be crys-stable. This is why we must introduce crys-stable bundles.

In order to keep track of how the bundle transforms after various applications of Frobenius, it is necessary to introduce a certain combinatorial device called a *VF-pattern* (where “VF” stands for “Verschiebung/Frobenius”). VF-patterns may be described as follows. Fix nonnegative integers g, r such that $2g - 2 + r > 0$. Let $\chi \stackrel{\text{def}}{=} \frac{1}{2}(2g - 2 + r)$. Let $\mathcal{L}ev$ be the set of $l \in \frac{1}{2}\mathbf{Z}$ satisfying $0 \leq l \leq \chi$. We shall call $\mathcal{L}ev$ the *set of levels*. (That is, $\mathcal{L}ev$ is the set of possible levels of crys-stable bundles.) Let $\Pi : \mathbf{Z} \rightarrow \mathcal{L}ev$ be a map of sets, and let ϖ be a positive integer. Then we make the following definitions:

(i) We shall call (Π, ϖ) a *VF-pattern* if $\Pi(n + \varpi) = \Pi(n)$ for all $n \in \mathbf{Z}$; $\Pi(0) = \chi$; $\Pi(i) - \Pi(j) \in \mathbf{Z}$ for all $i, j \in \mathbf{Z}$ (cf. Definition 1.1 of Chapter III).

(ii) A VF-pattern (Π, ϖ) will be called *pre-home* if $\Pi(\mathbf{Z}) = \{\chi\}$. A VF-pattern (Π, ϖ) will be called the *home VF-pattern* if it is pre-home and $\varpi = 1$.

(iii) A VF-pattern (Π, ϖ) will be called *binary* if $\Pi(\mathbf{Z}) \subseteq \{0, \chi\}$. A VF-pattern (Π, ϖ) will be called the *VF-pattern of pure tone ϖ* if $\Pi(n) = 0$ for all $n \in \mathbf{Z}$ not divisible by ϖ .

(iv) Let (Π, ϖ) be a VF-pattern. Then $i \in \mathbf{Z}$ will be called *indigenous* (respectively, *active*; *dormant*) for this VF-pattern if $\Pi(i) = \chi$ (respectively, $\Pi(i) \neq 0$; $\Pi(i) = 0$). If $i, j \in \mathbf{Z}$, and $i < j$, then (i, j) will be called *ind-adjacent*

for this VF-pattern if $\Pi(i) = \Pi(j) = \chi$ and $\Pi(n) \neq \chi$ for all $n \in \mathbf{Z}$ such that $i < n < j$.

At the present time, all of this terminology may seem rather abstruse, but eventually, we shall see that it corresponds in a natural and evident way to the p -adic geometry defined by indigenous bundles that are Frobenius invariant in the fashion described by the VF-pattern in question. Finally, we remark that often, in order to simplify notation, we shall just write Π for the VF-pattern (even though, strictly speaking, a VF-pattern is a pair (Π, ϖ)).

Now fix an odd prime p . Let (Π, ϖ) be a VF-pattern. Let S be a perfect scheme of characteristic p . Let $X \rightarrow S$ be a smooth, proper, geometrically connected curve of genus $g \geq 2$. (Naturally, the theory goes through for arbitrary pointed stable curves, but for simplicity, we assume in the present discussion that the curve is smooth and without marked points.) Write $W(S)$ for the (ind-)scheme of Witt vectors with coefficients in S . Let \mathcal{P} be a crystal on $\text{Crys}(X/W(S))$ valued in the category of \mathbf{P}^1 -bundles. Thus, the restriction $\mathcal{P}|_X$ of \mathcal{P} to $\text{Crys}(X/S)$ may be thought of as a \mathbf{P}^1 -bundle with connection on the curve $X \rightarrow S$. Let us assume that $\mathcal{P}|_X$ defines an *indigenous bundle* on X . Now we consider the following procedure (cf. Fig. 11):

Using the Hodge section of $\mathcal{P}|_X$, one can form the *renormalized Frobenius pull-back* $\mathcal{P}_1 \stackrel{\text{def}}{=} \mathbf{F}^*(\mathcal{P})$ of \mathcal{P} . Thus, $\mathbf{F}^*(\mathcal{P})$ will be a crystal in \mathbf{P}^1 -bundles on $\text{Crys}(X/W(S))$. Assume that $\mathcal{P}_1|_X$ is crys-stable of level $\Pi(1)$. Then there are two possibilities: either $\Pi(1)$ is zero or nonzero. If $\Pi(1) = 0$, then let \mathcal{P}_2 be the usual (i.e., non-renormalized) Frobenius pull-back $\Phi^*\mathcal{P}_1$ of the crystal \mathcal{P}_1 . If $\Pi(1) \neq 0$, then $\mathcal{P}_1|_X$ is crys-stable of positive level, hence admits a *Hodge section*; thus, using the Hodge section of $\mathcal{P}_1|_X$, we may form the renormalized Frobenius pull-back $\mathcal{P}_2 \stackrel{\text{def}}{=} \mathbf{F}^*(\mathcal{P}_1)$ of \mathcal{P}_1 . Continuing inductively in this fashion – i.e., always assuming $\mathcal{P}_i|_X$ to be crys-stable of level $\Pi(i)$, and forming \mathcal{P}_{i+1} by taking the renormalized (respectively, usual) Frobenius pull-back of \mathcal{P}_i if $\Pi(i) \neq 0$ (respectively, $\Pi(i) = 0$), we obtain a sequence \mathcal{P}_i of crystals on $\text{Crys}(X/W(S))$ valued in the category of \mathbf{P}^1 -bundles.

Then we make the following

Definition 1.5. We shall refer to \mathcal{P} as Π -indigenous (on X) if all the assumptions (on the \mathcal{P}_i) necessary to carry out the above procedure are satisfied, and, moreover, $\mathcal{P}_\varpi \cong \mathcal{P}$.

Thus, to say that \mathcal{P} is Π -indigenous (more properly, (Π, ϖ) -indigenous) is to say that it is Frobenius invariant in the fashion specified by the combinatorial data (Π, ϖ) .

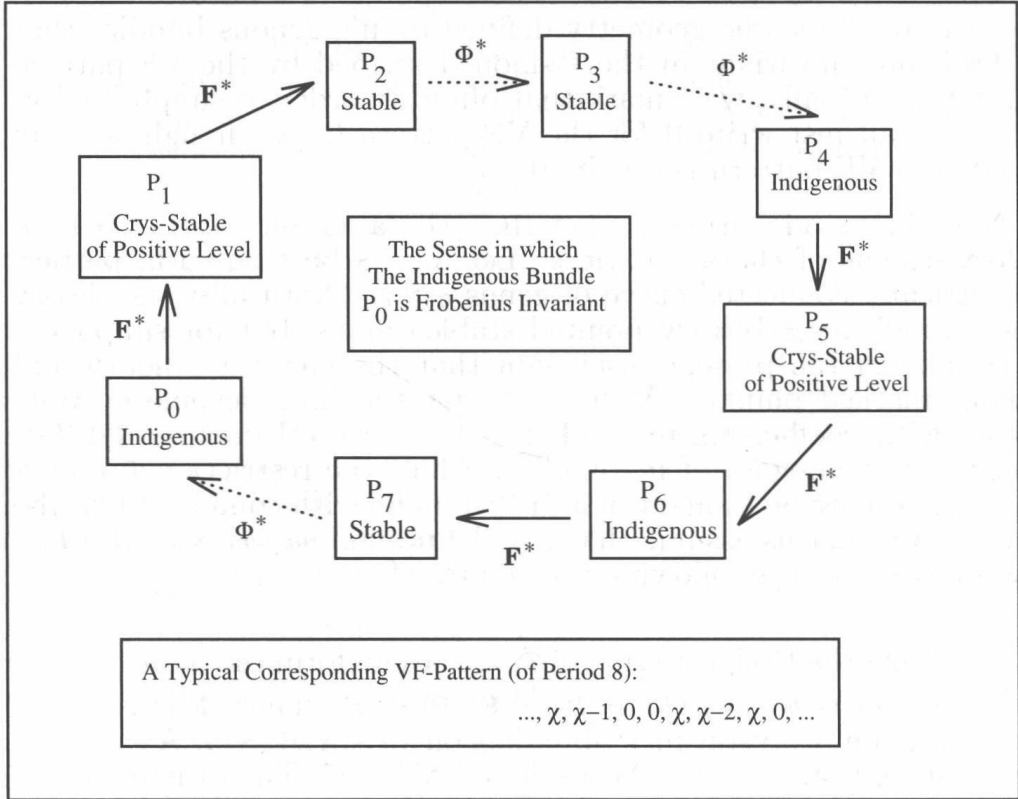


Fig. 11: The Sense of Frobenius Invariance Specified by a VF-Pattern

Now we are ready to define a certain stack that will be of central importance in this work. The stack \mathcal{Q}^Π – also called the *stack of quasi-analytic self-isogenies of type (Π, ϖ)* – is defined as follows:

To a perfect scheme S , $\mathcal{Q}^\Pi(S)$ assigns the category of pairs $(X \rightarrow S, \mathcal{P})$, where $X \rightarrow S$ is a curve as above and \mathcal{P} is a Π -indigenous bundle on X .

Thus, \mathcal{Q}^Π may be regarded as the moduli stack of indigenous bundles that are Frobenius invariant in the fashion specified by the VF-pattern Π .

We remark that in fact, more generally, one can define \mathcal{Q}^Π on the category of *epiperfect schemes* S . (Whereas a perfect scheme is a scheme on which the Frobenius morphism is an isomorphism, an epiperfect scheme is one on which the Frobenius morphism is a closed immersion.) Then instead of using $W(S)$, one works over $B(S)$ – i.e., the

“universal PD -thickening of S .” For instance, the well-known ring B_{crys} of Fontaine is a special case of $B(S)$. The point is that one needs the base spaces that one works with to be \mathbf{Z}_p -flat and equipped with a natural Frobenius action. The advantage of working with arbitrary $B(S)$ (for S epiperfect) is that the theory of crystalline representations (and the fact that B_{crys} is a special case of $B(S)$) suggest that $B(S)$ is likely to be the most general natural type of space having these two properties – i.e., \mathbf{Z}_p -flatness and being equipped with a natural Frobenius action. The disadvantage of working with arbitrary $B(S)$ (as opposed to just $W(S)$ for perfect S) is that many properties of \mathcal{Q}^Π are technically more difficult or (at the present time impossible) to prove in the epiperfect case. For the sake of simplicity, in this Introduction, we shall only consider the perfect case. For more details, we refer to Chapter VI.

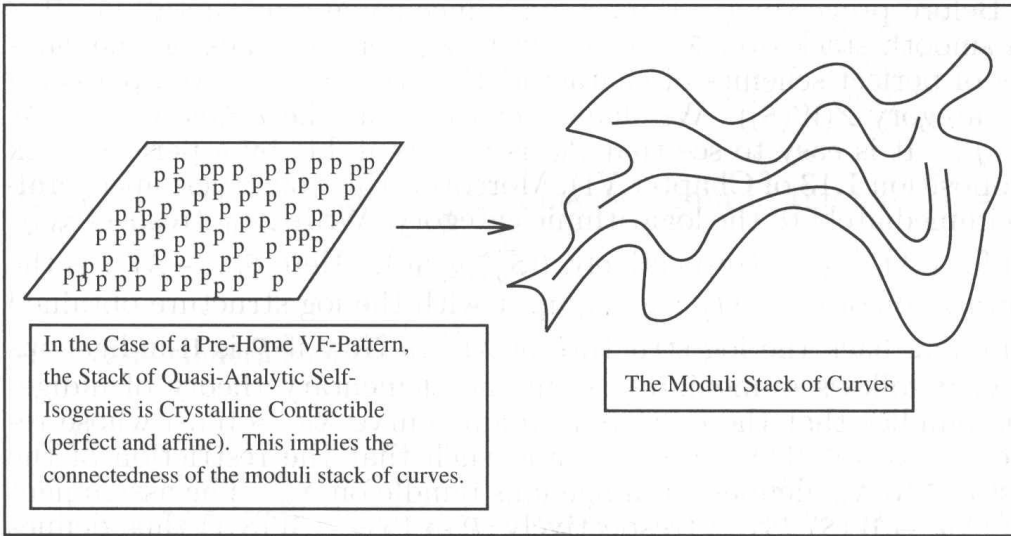


Fig. 12: Crystalline Contractibility in the Pre-Home Case

Now, we are ready to discuss the main results concerning \mathcal{Q}^Π . The general theory of \mathcal{Q}^Π is the topic of Chapter VI. We begin with the following result (cf. Theorem 2.2 of Chapter VI):

Theorem 1.6. (Representability and Affineness) *The stack \mathcal{Q}^Π is representable by a perfect algebraic stack whose associated coarse moduli space (as in [FC], Chapter 1, Theorem 4.10) is quasi-affine. If Π is pre-home, then this coarse moduli space is even affine.*

Thus, in the pre-home case, \mathcal{Q}^Π is *perfect and affine*. In particular, any sort of de Rham/crystalline-type cohomology on \mathcal{Q}^Π must vanish. It is for this reason that we say (in the pre-home case) that \mathcal{Q}^Π is *crystalline*

contractible (cf. Fig. 12). Moreover, as discussed earlier in §1.1 (cf. Theorem 2.12 of Chapter III),

Corollary 1.7. (Irreducibility of Moduli) *(The fact that \mathcal{Q}^Π is crystalline contractible for the home VF-pattern is intimately connected with the fact that $\mathcal{M}_{g,r}$ is irreducible.*

As remarked earlier, this is reminiscent of the proof of the irreducibility of $\mathcal{M}_{g,r}$ given by using complex Teichmüller theory to show that Teichmüller space is contractible. Moreover, it is also interesting in that it suggests that perhaps at some future date the theory (or some extension of the theory) of the present work may be used to compute other cohomology groups of $\mathcal{M}_{g,r}$.

Before proceeding, we must introduce some more notation. If Z is a smooth stack over \mathbf{Z}_p , let us write Z_W for the stack on the category of perfect schemes of characteristic p that assigns to a perfect S the category $Z(W(S))$. We shall refer to Z_W as the *infinite Weil restriction* of Z . It is easy to see that Z_W is representable by a perfect stack (Proposition 1.13 of Chapter VI). Moreover, this construction generalizes immediately to the logarithmic category. Write \mathcal{M}_W (respectively, \mathcal{S}_W) for $((\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{Z}_p})_W$ (respectively, $((\overline{\mathcal{S}}_{g,r}^{\log})_{\mathbf{Z}_p})_W$). (Here $\overline{\mathcal{S}}_{g,r} \rightarrow \overline{\mathcal{M}}_{g,r}$ is the Schwarz torsor over $\overline{\mathcal{M}}_{g,r}$; we equip it with the log structure obtained by pulling back the log structure of $\overline{\mathcal{M}}_{g,r}^{\log}$.) Now if \mathcal{P} is Π -indigenous on X , it follows immediately from the elementary theory of indigenous bundles that there exists a unique curve $X_W \rightarrow W(S)$ whose restriction to $S \subseteq W(S)$ is $X \rightarrow S$ and such that the restriction of the crystal \mathcal{P} to X_W defines an indigenous bundle on X_W . The assignment $\mathcal{P} \mapsto (X_W \rightarrow W(S), \mathcal{P}|_{X_W})$ (respectively, $\mathcal{P} \mapsto \{X_W \rightarrow W(S)\}$) thus defines a natural morphism $\mathcal{Q}^\Pi \rightarrow \mathcal{S}_W$ (respectively, $\mathcal{Q}^\Pi \rightarrow \mathcal{M}_W$). Now we have the following results (cf. Propositions 2.3, 2.9; Corollaries 2.6 and 2.13 of Chapter VI):

Theorem 1.8. (Immersion) *The natural morphism $\mathcal{Q}^\Pi \rightarrow \mathcal{S}_W$ is an immersion in general, and a closed immersion if the VF-pattern is pre-home or of pure tone. The morphism $\mathcal{Q}^\Pi \rightarrow \mathcal{M}_W$ is a closed immersion if the VF-pattern is the home VF-pattern.*

Theorem 1.9. (Isolatedness in the Pre-Home Case) *In the pre-home case, \mathcal{Q}^Π is closed inside \mathcal{S}_W and disjoint from the closure of any non-pre-home $\mathcal{Q}^{\Pi'}$'s.*

We remark that in both of these cases, much more general theorems are proved in the text. Here, for the sake of simplicity, we just selected representative special cases of the main theorems in the text so as to

give the reader a general sense of the sorts of results proved in the text.

The reason that Theorem 1.9 is interesting (or perhaps a bit surprising) is the following: The reduction modulo p of a Π -indigenous bundle (in the pre-home case) is an *admissible* nilpotent indigenous bundle. (Recall that “admissible” means that the p -curvature has no zeroes.) Moreover, the admissible locus $\overline{\mathcal{N}}_{g,r}^{\text{adm}}$ of $\overline{\mathcal{N}}_{g,r}$ is by no means closed in $\overline{\mathcal{N}}_{g,r}$, nor is its closure disjoint (in general) from the closure of the dormant or spiked loci of $\overline{\mathcal{N}}_{g,r}$. On the other hand, the reductions modulo p of Π' -indigenous bundles (for non-pre-home Π') may, in general, be dormant or spiked nilpotent indigenous bundles. Thus,

Theorem 1.9 states that considering \mathbf{Z}_p -flat Frobenius invariant liftings of indigenous bundles (as opposed to just nilpotent indigenous bundles in characteristic p) has the effect of “blowing up” $\overline{\mathcal{N}}_{g,r}$ in such a way that the genericization/specialization relations that hold in $\overline{\mathcal{N}}_{g,r}$ do not imply such relations among the various \mathcal{Q} ’s.

We shall come back to this phenomenon again in the following subsection (cf. Fig. 13).

§1.5. The Generalized Ordinary Theory

In this subsection, we maintain the notations of the preceding subsection. Unfortunately, it is difficult to say much more about the explicit structure of the stacks \mathcal{Q}^Π without making more assumptions. Thus, just as in the classical ordinary case (reviewed in §0.9), it is natural to define an open substack – the *ordinary locus* of \mathcal{Q}^Π – and to see if more explicit things can be said concerning this open substack. This is the topic of Chapter VII. We shall see below that in fact much that is interesting can be said concerning this ordinary locus.

We begin with the definition of the ordinary locus. First of all, we observe that there is a natural algebraic stack

$$\overline{\mathcal{N}}_{g,r}^{\Pi,s}$$

(of finite type over \mathbf{F}_p) that parametrizes “data modulo p for \mathcal{Q}^Π ” (Definition 1.11 of Chapter III). That is to say, roughly speaking, $\overline{\mathcal{N}}_{g,r}^{\Pi,s}$ parametrizes the reductions modulo p of the \mathcal{P}_i appearing in the discussion preceding Definition 1.5. We refer to Chapter III for a precise definition of this stack. At any rate, by reducing modulo p the data parametrized by \mathcal{Q}^Π , we obtain a natural morphism of stacks

$$\mathcal{Q}^\Pi \rightarrow \overline{\mathcal{N}}_{g,r}^{\Pi,s}$$

On the other hand, since $\overline{\mathcal{N}}_{g,r}^{\Pi,s}$ parametrizes curves equipped with certain bundles, there is a natural morphism $\overline{\mathcal{N}}_{g,r}^{\Pi,s} \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$. Let $\mathcal{N}^{\text{ord}} \subseteq \overline{\mathcal{N}}_{g,r}^{\Pi,s}$ denote the open substack over which the morphism $\overline{\mathcal{N}}_{g,r}^{\Pi,s} \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ is *étale*. Let $\mathcal{Q}^{\text{ord}} \subseteq \mathcal{Q}^\Pi$ denote the open substack which is the inverse image of $\mathcal{N}^{\text{ord}} \subseteq \overline{\mathcal{N}}_{g,r}^{\Pi,s}$.

Definition 1.10. We shall refer to \mathcal{Q}^{ord} as the $(\Pi\text{-})$ *ordinary locus* of \mathcal{Q}^Π .

Just as in the classical ordinary case, there is an equivalent definition of Π -ordinariness given by looking at the action of Frobenius on the first de Rham cohomology modules of the \mathcal{P}_i (cf. Lemma 1.4 of Chapter VII). Incidentally, the classical ordinary theory corresponds to the Π -ordinary theory in the case of the home VF-pattern. (In particular, \mathcal{N}^{ord} is simply the ordinary locus $\overline{\mathcal{N}}_{g,r}^{\text{ord}}$ of $\overline{\mathcal{N}}_{g,r}$.) Thus, in some sense, the theory of [Mzk1] is a special case of the generalized ordinary theory.

Our first result is the following (cf. Theorem 1.6 of Chapter VII):

Theorem 1.11. (Basic Structure of the Ordinary Locus) \mathcal{Q}^{ord} is naturally isomorphic to the perfection of \mathcal{N}^{ord} .

Thus, already one has a much more explicit understanding of the structure of \mathcal{Q}^{ord} than of the whole of \mathcal{Q}^Π . That is to say, Theorem 1.11 already tells us that \mathcal{Q}^{ord} is the perfection of a smooth algebraic stack of finite type over \mathbf{F}_p .

Our next result – which is somewhat deeper than Theorem 1.11, and is, in fact, one of the main results of this book – is the following (cf. Theorem 2.11 of Chapter VII):

Theorem 1.12. (ω -Closedness of the Ordinary Locus) If Π is binary, then \mathcal{Q}^{ord} is ω -closed (roughly speaking, “closed as far as the differentials are concerned” – cf. Chapter VII, §0, §2.3 for more details) in \mathcal{Q}^Π . In particular,

- (1) If $3g - 3 + r = 1$, then \mathcal{Q}^{ord} is actually closed in \mathcal{Q}^Π .
- (2) If $\mathcal{R} \subseteq \mathcal{Q}^\Pi$ is a subobject containing \mathcal{Q}^{ord} and which is “pro” (cf. Chapter VI, Definition 1.9) of a fine algebraic log stack which is locally of finite type over \mathbf{F}_p , then \mathcal{Q}^{ord} is closed in \mathcal{R} .

In other words, at least among perfection of fine algebraic log stacks which are locally of finite type over \mathbf{F}_p , \mathcal{Q}^{ord} is already “complete” inside \mathcal{Q}^{Π} .

Thus, if Π is pre-home or of pure tone, then \mathcal{Q}^{ord} is an ω -closed substack of S_W . If the VF-pattern in question is the home pattern, then \mathcal{Q}^{ord} is an ω -closed substack of \mathcal{M}_W .

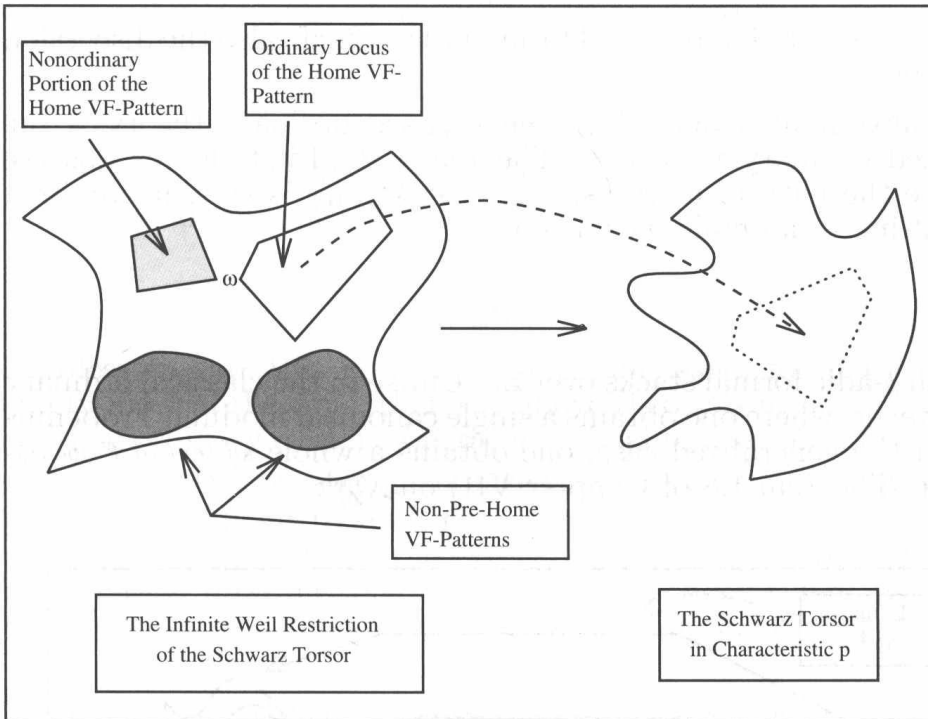


Fig. 13: The ω -Closedness and Isolatedness of the Classical Ordinary Theory

This is a rather surprising result in that the definition of \mathcal{Q}^{ord} was such that \mathcal{Q}^{ord} is naturally an open substack of \mathcal{Q}^{Π} which has no a priori reason to be closed (in any sense!) inside \mathcal{Q}^{Π} . Moreover, $\overline{\mathcal{N}}_{g,r}^{\text{ord}}$ is most definitely *not* closed in $\overline{\mathcal{N}}_{g,r}$. Indeed, one of the original motivations for trying to generalize the theory of [Mzk1] was to try to compactify it. Thus, Theorem 1.12 states that if, instead of just considering ordinary nilpotent indigenous bundles modulo p , one considers \mathbf{Z}_p -flat Frobenius invariant indigenous bundles, the theory of [Mzk1] is, in some sense, already compact! Put another way, if one thinks in terms of the Witt vectors parametrizing such \mathbf{Z}_p -flat Frobenius invariant indigenous bundles, then although the scheme defined by the *first* component of the Witt vector is not “compact,” if one considers *all* the components of the Witt vector, the resulting scheme is, in

some sense, “compact” (i.e., ω -closed in the space \mathcal{S}_W of all indigenous bundles over the Witt vectors). This phenomenon is similar to the phenomenon observed in Theorem 1.9. In fact, if one combines Theorem 1.9 with Theorem 1.12, one obtains that:

In the home (i.e., classical ordinary) case, the stack \mathcal{Q}^{ord} is ω -closed in \mathcal{S}_W and disjoint from the closures of all $\mathcal{Q}^{\Pi'}$ for all non-pre-home Π' . Moreover, \mathcal{Q}^{ord} is naturally an ω -closed substack of $\mathcal{Q}^{\Pi'}$ for all pre-home Π' .

This fact is rendered in pictorial form in Fig. 13; cf. also the discussion of Chapter X, §3.

The next main result of the generalized ordinary theory is the generalized ordinary version of Theorem 0.2. First, let us observe that since the natural morphism $\mathcal{N}^{\text{ord}} \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbb{F}_p}$ is étale, it admits a unique lifting to an étale morphism

$$\mathcal{N}_{\mathbb{Z}_p}^{\text{ord}} \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbb{Z}_p}$$

of smooth p -adic formal stacks over \mathbb{Z}_p . Unlike in the classical ordinary case, however, where one obtains a single canonical modular Frobenius lifting, in the generalized case, one obtains a whole *system of Frobenius liftings* (cf. Theorem 1.8 of Chapter VII) on $\mathcal{N}_{\mathbb{Z}_p}^{\text{ord}}$:




Fig. 14: The Canonical System of Modular Frobenius Liftings

Theorem 1.13. (**Canonical System of Frobenius Liftings**) *Over $\mathcal{N}_{\mathbf{Z}_p}^{\text{ord}}$, there is a canonical system of Frobenius liftings and indigenous bundles: i.e., for each indigenous i (i.e., such that $\Pi(i) = \chi$), a lifting*

$$\Phi_i^{\log} : \mathcal{N}_{\mathbf{Z}_p}^{\text{ord}} \rightarrow \mathcal{N}_{\mathbf{Z}_p}^{\text{ord}}$$

of a certain power of the Frobenius morphism, together with a collection of indigenous bundles \mathcal{P}_i on the tautological curve (pulled back from $(\mathcal{M}_{g,r})_{\mathbf{Z}_p}$) over $\mathcal{N}_{\mathbf{Z}_p}^{\text{ord}}$. Moreover, these Frobenius liftings and bundles are compatible, in a natural sense (Definition 1.7 of Chapter VII).

See Fig. 14 for an illustration of the system of Frobenius liftings obtained for the VF-pattern illustrated in Fig. 11.

At this point, one very important question arises:

To what extent are the stacks \mathcal{N}^{ord} nonempty?

Needless to say, this is a very important issue, for if the \mathcal{N}^{ord} are empty most of the time, then the above theory is meaningless. In the classical ordinary case, it was rather trivial to show the nonemptiness of $\overline{\mathcal{N}}_{g,r}^{\text{ord}}$. In the present generalized ordinary setting, however, it is much more difficult to show the nonemptiness of \mathcal{N}^{ord} . In particular, one needs to make use of the extensive theory of Chapters II and IV. Fortunately, however, one can show the nonemptiness of \mathcal{N}^{ord} in a fairly wide variety of cases (Theorems 3.1 and 3.7 of Chapter VII):

Theorem 1.14. (**Binary Existence Result**) *Suppose that $g \geq 2$; $r = 0$; and $p > 4^{3g-3}$. Then for any binary VF-pattern (i.e., VF-pattern such that $\Pi(\mathbf{Z}) \subseteq \{0, \chi\}$), the stack \mathcal{N}^{ord} will be nonempty.*

Theorem 1.15. (**Spiked Existence Result**) *Suppose that $2g - 2 + r \geq 3$ and $p \geq 5$. Then there exists a “spiked VF-pattern” of period 2 (i.e., $\varpi = 2$ and $0 < \Pi(1) < \chi$) for which \mathcal{N}^{ord} is nonempty.*

In fact, there is an open substack of \mathcal{N}^{ord} called the very ordinary locus (defined by more stringent conditions than ordinarity); moreover, one can choose the spiked VF-pattern so that not only \mathcal{N}^{ord} , but also the “very ordinary locus of \mathcal{N}^{ord} ” is nonempty.

These cases are “fairly representative” in the following sense: In general, in the binary case, the reduction modulo p of a Π -indigenous

bundle will be *dormant*. In the spiked case (of Theorem 1.15), the reduction modulo p of a Π -indigenous bundle will be *spiked*. Thus, in other words,

Roughly speaking, these two existence results show that for each type (admissible, dormant, spiked) of nilcurve, there exists a theory (in fact, many theories) of canonical liftings involving that type of nilcurve.

Showing the existence of such a theory of canonical liftings for *each generic point* of $\overline{\mathcal{N}}_{g,r}$ was one of the original motivations for the development of the theory of this book.

Next, we observe that just as in Theorem 0.2 (the classical ordinary case),

In the cases discussed in Theorems 1.14 and 1.15, one can also construct canonical systems of Frobenius liftings on certain “ordinary loci” of the universal curve over $\mathcal{N}_{\mathbf{Z}_p}^{\text{ord}}$. Moreover, these systems of canonical Frobenius lifting lie over the canonical system of modular Frobenius liftings of Theorem 1.13.

We refer to Theorem 3.2 of Chapter VIII and Theorem 3.4 of Chapter IX for more details.

We end this subsection with a certain philosophical observation. In Chapter VI,

The stack \mathcal{Q}^Π is referred to as the stack of quasi-analytic self-isogenies.

That is to say, in some sense it is natural to regard the Frobenius invariant indigenous bundles parametrized by \mathcal{Q}^Π as isogenies of the curve (on which the bundles are defined) onto itself. Indeed, this is suggested by the fact that over the ordinary locus (i.e., relative to the Frobenius invariant indigenous bundle in question) of the curve, the bundle actually does define a literal morphism, i.e., a Frobenius lifting (as discussed in the preceding paragraph). Thus, one may regard a Frobenius invariant indigenous bundle as the appropriate way of compactifying such a self-isogeny to an object defined over the whole curve. This is why we use the adjective “quasi-analytic” in describing the self-isogenies. (Of course, such self-isogenies can never be p -adic analytic over the whole curve, for if they were, they would be algebraic, which, by the Riemann-Hurwitz formula, is absurd.) Note that this point of view is in harmony with the situation in the parabolic case ($g = 1, r = 0$), where there is an algebraically defined canonical choice of indigenous bundle, and having a Frobenius invariant indigenous bundle really does correspond to having a lifting of Frobenius (hence a self-isogeny of the curve in question).

Moreover, note that in the case where the VF-pattern has several $\chi = \frac{1}{2}(2g - 2 + r)$'s in a period, so that there are various indigenous \mathcal{P}_i 's in addition to the original Frobenius invariant indigenous bundle \mathcal{P} , one may regard the situation as follows. Suppose that \mathcal{P} is indigenous over a curve $X \rightarrow W(S)$, whereas \mathcal{P}_i is indigenous over $X_i \rightarrow W(S)$. Then one can regard the “quasi-analytic self-isogeny” $\mathcal{P} : X \rightarrow X$ as the composite of various quasi-analytic isogenies $\mathcal{P}_i : X_i \rightarrow X_j$ (where i and j are “ind-adjacent” integers). Note that this point of view is consistent with what literally occurs over the ordinary locus (cf. Theorem 3.2 of Chapter VIII). Finally, we observe that

The idea that \mathcal{Q}^Π is a moduli space of some sort of p-adic self-isogeny which is “quasi-analytic” is also compatible with the analogy between \mathcal{Q}^Π and Teichmüller space (cf. the discussion of Corollary 1.7) in that Teichmüller space may be regarded as a moduli space of quasiconformal maps.

§1.6. Geometrization

In the classical ordinary case, once one knows the existence of the canonical modular Frobenius lifting (Theorem 0.2), one can apply a general result on ordinary Frobenius liftings (Theorem 0.3) to conclude the existence of canonical multiplicative coordinates on $\mathcal{N}_{\mathbb{Z}_p}^{\text{ord}}$. We shall refer to this process of passing (as in Theorem 0.3) from a certain type of Frobenius lifting to a local uniformization/canonical local coordinates associated to the Frobenius lifting as the *geometrization of the Frobenius lifting*. In the generalized ordinary context, Theorem 1.13 shows the existence of a canonical *system of Frobenius liftings* on the $\mathcal{N}_{\mathbb{Z}_p}^{\text{ord}}$ associated to a VF-pattern (Π, ϖ) . Thus, the following question naturally arises:

Can one geometrize the sort of system of Frobenius liftings that one obtains in Theorem 1.13 in a fashion analogous to the way in which ordinary Frobenius liftings were geometrized in Theorem 0.3?

Unfortunately, we are not able to answer this question in general. Nevertheless, in the cases discussed in Theorems 1.14 and 1.15, i.e., the binary and very ordinary spiked cases, we succeed (in Chapters VIII and IX) in geometrizing the canonical system of modular Frobenius liftings. The result is *uniformizations/geometries based not on $\hat{\mathbf{G}}_m$ as in the classical ordinary case, but rather on more general types of Lubin-Tate groups, twisted products of Lubin-Tate groups, and fibrations whose bases are Lubin-Tate groups and whose fibers are such twisted products*. In the rest of this subsection, we would like to try to give the reader an idea of what sorts of geometries occur in the two cases studied.

In the following, we let k be a perfect field of characteristic p , A its ring of Witt vectors $W(k)$, and S a smooth p -adic formal scheme over A . Let λ be a positive integer, and let $\mathcal{O}_\lambda \stackrel{\text{def}}{=} W(\mathbf{F}_{p^\lambda})$. For simplicity, we assume that $\mathcal{O}_\lambda \subseteq A$. Let \mathcal{G}_λ be the *Lubin-Tate formal group associated to \mathcal{O}_λ* . (See [CF] for a discussion of Lubin-Tate formal groups.) Then \mathcal{G}_λ is a formal group over \mathcal{O}_λ , equipped with a natural action by \mathcal{O}_λ (i.e., a ring morphism $\mathcal{O}_\lambda \hookrightarrow \text{End}_{\mathcal{O}_\lambda}(\mathcal{G}_\lambda)$). Moreover, it is known that the space of invariant differentials on \mathcal{G}_λ is canonically isomorphic to \mathcal{O}_λ . Thus, in the following, we shall identify this space of differentials with \mathcal{O}_λ .

We begin with the simplest case, namely, that of a *Lubin-Tate Frobenius lifting*. Let $\Phi : S \rightarrow S$ be a morphism whose reduction modulo p is the λ^{th} power of the Frobenius morphism. Then differentiating Φ_S defines a morphism $d\Phi_S : \Phi_S^* \Omega_{S/A} \rightarrow \Omega_{S/A}$ which is zero in characteristic p . Thus, we may form a morphism

$$\Omega_\Phi : \Phi_S^* \Omega_{S/A} \rightarrow \Omega_{S/A}$$

by dividing $d\Phi_S$ by p . Then Φ_S is called a *Lubin-Tate Frobenius lifting (of order λ)* if Ω_Φ is an isomorphism. If Φ_S is a Lubin-Tate Frobenius lifting, then it induces a “Lubin-Tate geometry” – i.e., a geometry based on \mathcal{G}_λ – on S . That is to say, one has the following analogue of Theorem 0.3 (cf. Theorem 2.17 of Chapter VIII):

Theorem 1.16. (Lubin-Tate Frobenius Liftings) *Let $\Phi_S : S \rightarrow S$ be a Lubin-Tate Frobenius lifting of order λ . Then taking the invariants of $\Omega_{S/A}$ with respect to Ω_Φ gives rise to an étale local system $\Omega_\Phi^{\text{ét}}$ on S of free \mathcal{O}_λ -modules of rank equal to $\dim_A(S)$.*

Let $z \in S(\bar{k})$ be a point valued in the algebraic closure of k . Then $\Omega_z \stackrel{\text{def}}{=} \Omega_\Phi^{\text{ét}}|_z$ may be thought of as a free \mathcal{O}_λ -module of rank $\dim_A(S)$; write Θ_z for the \mathcal{O}_λ -dual of Ω_z . Let S_z be the completion of S at z . Then there is a unique isomorphism

$$\Gamma_z : S_z \cong \mathcal{G}_\lambda \otimes_{\mathcal{O}_\lambda}^{\text{gp}} \Theta_z$$

such that:

- (i) *the derivative of Γ_z induces the natural inclusion $\Omega_z \hookrightarrow \Omega_{S/A}|_{S_z}$;*
- (ii) *the action of Φ_S on S_z corresponds to multiplication by p on $\mathcal{G}_\lambda \otimes_{\mathcal{O}_\lambda}^{\text{gp}} \Theta_z$.*

Here, by “ $\mathcal{G}_\lambda \otimes_{\mathcal{O}_\lambda}^{\text{gp}} \Theta_z$,” we mean the tensor product over \mathcal{O}_λ of (formal) group schemes with \mathcal{O}_λ -action. Thus, $\mathcal{G}_\lambda \otimes_{\mathcal{O}_\lambda}^{\text{gp}} \Theta_z$ is noncanonically isomorphic to the product of $\dim_A(S) = \text{rank}_{\mathcal{O}_\lambda}(\Theta_z)$ copies of \mathcal{G}_λ .

Of course, this result has nothing to do with the moduli of curves. In terms of VF-patterns, Theorem 1.13 gives rise to a Lubin-Tate Frobenius lifting of order ϖ when the VF-pattern is of pure tone ϖ .

The next simplest case is the case of an *anabelian system of Frobenius liftings*. Let n be a positive integer. Then an *anabelian system of Frobenius liftings of length n and order λ* is a collection of n Lubin-Tate Frobenius liftings

$$\Phi_1, \dots, \Phi_n : S \rightarrow S$$

each of order λ . Of course, in general such Frobenius liftings will not commute with one another. In fact, it can be shown that two Lubin-Tate Frobenius liftings of order λ commute with each other if and only if they are equal (Lemma 2.24 of Chapter VIII). This is the reason for the term “anabelian.” Historically, this term has been used mainly in connection with the Grothendieck Conjecture (cf. the discussion of §0.10). The reason why we thought it appropriate to use the term here (despite the fact that anabelian geometries as discussed here have nothing to do with the Grothendieck Conjecture) is the following: *The sort of noncommutativity that occurs among the Φ_i ’s (at least in the modular case discussed in Theorem 1.13) arises precisely as a result of the hyperbolicity of the curves on whose moduli the Φ_i ’s act.* That is to say, for instance, for elliptic curves regarded *parabolically* (i.e., the case $g = 1, r = 0$), such anabelian geometries do not arise precisely because $N_{1,0}^{\text{ord}}$ is of degree 1 over $(\mathcal{M}_{1,0})_{\mathbb{F}_p}$, whereas for elliptic curves regarded *hyperbolically* (i.e., the case $g = 1, r = 1$), such anabelian geometries arise precisely because $N_{1,1}^{\text{ord}}$ is of degree > 1 (in general) over $(\mathcal{M}_{1,1})_{\mathbb{F}_p}$. In other words,

In both the present case of “anabelian geometries” and the case of the Grothendieck Conjecture, the term anabelian refers not just to “noncommutativity” but specifically the sort of noncommutativity that arises precisely as a result of hyperbolicity.

We shall say more about the noncommutative nature of anabelian geometries after we state the main theorem (Theorem 1.17 below) on these geometries.

Let $\delta_i \stackrel{\text{def}}{=} \frac{1}{p} d\Phi_i$. Let $\Delta \stackrel{\text{def}}{=} \delta_n \circ \dots \circ \delta_1$. Then taking invariants of $\Omega_{S/A}$ with respect to Δ gives rise to an étale local system $\Omega_{\Phi}^{\text{et}}$ on S in free $\mathcal{O}_{n\lambda}$ -modules of rank $\dim_A(S)$. Next let S_{PD} denote the p -adic completion of the PD -envelope of the diagonal in the product (over A) of n copies of S ; let S_{FM} denote the p -adic completion of the completion at the diagonal of the product (over A) of n copies of S . Thus, we have a natural morphism

$$S_{\text{PD}} \rightarrow S_{\text{FM}}$$

Moreover, one may think of S_{PD} as a sort of *localization* of S_{FM} . Write $\Phi_{\text{PD}} : S_{\text{PD}} \rightarrow S_{\text{PD}}$ for the morphism induced by sending

$$(s_1, \dots, s_n) \mapsto (\Phi_1(s_2), \Phi_2(s_3), \dots, \Phi_n(s_1))$$

(where (s_1, \dots, s_n) represents a point in the product of n copies of S). Then we have the following result (cf. Theorem 2.17 of Chapter VIII):

Theorem 1.17. (Anabelian System of Frobenius Liftings) *Let $\Phi_1, \dots, \Phi_n : S \rightarrow S$ be a system of anabelian Frobenius liftings of length n and order λ . Let $z \in S(\bar{k})$ be a point valued in the algebraic closure of k . Then $\Omega_z \stackrel{\text{def}}{=} \Omega_{\Phi}^{\text{et}}|_z$ may be thought of as a free $\mathcal{O}_{n\lambda}$ -module of rank $\dim_A(S)$; write Θ_z for the $\mathcal{O}_{n\lambda}$ -dual of Ω_z . Let $(S_{\text{PD}})_z$ be the completion of S_{PD} at z . Then there is a unique morphism*

$$\Gamma_z : (S_{\text{PD}})_z \rightarrow \mathcal{G}_{\lambda} \otimes_{\mathcal{O}_{\lambda}}^{\text{gp}} \Theta_z$$

such that:

(i) *the derivative of Γ_z induces a certain (see Theorem 2.15 of Chapter VIII for more details) natural inclusion of Ω_z into the restriction to $(S_{\text{PD}})_z$ of the differentials of $\prod_{i=1}^n S$ over A ;*

(ii) *the action of Φ_{PD} on $(S_{\text{PD}})_z$ is compatible with multiplication by p on $\mathcal{G}_{\lambda} \otimes_{\mathcal{O}_{\lambda}}^{\text{gp}} \Theta_z$.*

Here, by “ $\mathcal{G}_{\lambda} \otimes_{\mathcal{O}_{\lambda}}^{\text{gp}} \Theta_z$,” we mean the tensor product over \mathcal{O}_{λ} of (formal) group schemes with \mathcal{O}_{λ} -action. Thus, $\mathcal{G}_{\lambda} \otimes_{\mathcal{O}_{\lambda}}^{\text{gp}} \Theta_z$ is noncanonically isomorphic to the product of $n \cdot \dim_A(S) = \text{rank}_{\mathcal{O}_{\lambda}}(\Theta_z)$ copies of \mathcal{G}_{λ} .

Moreover, in general, Γ_z does not descend to $(S_{\text{FM}})_z$ (cf. Chapter VIII, §2.6, 3.1).

One way to envision anabelian geometries is as follows: The various Φ_i ’s induce various linear Lubin-Tate geometries on the space S that (in general) do not commute with one another. Thus, *the anabelian geometry consists of various linear geometries on S all tangled up inside each other*. If one localizes in a sufficiently drastic fashion – i.e., all the way to $(S_{\text{PD}})_z$ – then one can untangle these tangled up linear geometries into a single $\mathcal{O}_{n\lambda}$ -linear geometry (via Γ_z). However, the order λ Lubin-Tate geometries are so tangled up that even localization to a relatively localized object such as $(S_{\text{FM}})_z$ is not sufficient to untangle these geometries.

Finally, to make the connection with Theorem 1.13, we remark that the system of Theorem 1.13 gives rise to an anabelian system of length n and order λ in the case of a VF-pattern (Π, ϖ) for which

$\varpi = n \cdot \lambda$, and $\Pi(i) = \chi$ (respectively, $\Pi(i) = 0$) if and only if i is divisible (respectively, not divisible) by λ .

In fact, both Lubin-Tate geometries and anabelian geometries are special cases of *binary ordinary geometries* (the sorts of geometries that occur for *binary VF-patterns*, i.e., Π whose image $\subseteq \{0, \chi\}$). A general geometrization result for binary ordinary geometries is given in Theorem 2.17 of Chapter VIII. Here, we chose to concentrate on the Lubin-Tate and anabelian cases (in fact, of course, Lubin-Tate geometries are a special case of anabelian geometries) since they are relatively representative and relatively easy to envision.

The other main type of geometry that is studied in this book is the geometry associated to a *very ordinary spiked Frobenius lifting* $\Phi : S \rightarrow S$. Such a Frobenius lifting reduces modulo p to the square of the Frobenius morphism and satisfies various other properties which we omit here (see Definition 1.1 of Chapter IX for more details). In particular, such a Frobenius lifting comes equipped with an invariant called the *colevel*. The colevel is a nonnegative integer c . Roughly speaking,

A very ordinary spiked Frobenius lifting is a Frobenius lifting which is “part Lubin-Tate of order 2” and “part anabelian of length 2 and order 1.”

The colevel c is the number of dimensions of S on which Φ is Lubin-Tate of order 2. The main geometrization theorem (roughly stated) on this sort of Frobenius lifting is as follows (cf. Theorems 1.5 and 2.3 of Chapter IX):

Theorem 1.18. (Very Ordinary Spiked Frobenius Liftings) *Let $\Phi : S \rightarrow S$ be a very ordinary spiked Frobenius lifting of colevel c . Then Φ defines an étale local system Ω_Φ^{st} on S of free \mathcal{O}_2 -modules of rank c equipped with a natural inclusion $\Omega_\Phi^{\text{st}} \hookrightarrow \Omega_{S/A}$.*

Let $z \in S(\bar{k})$ be a point valued in the algebraic closure of k . Then $\Omega_z^{\text{st}} \stackrel{\text{def}}{=} \Omega_\Phi^{\text{st}}|_z$ may be thought of as a free \mathcal{O}_2 -module of rank c ; write Θ_z^{st} for the \mathcal{O}_2 -dual of Ω_z^{st} . Let S_z be the completion of S at z . Then there is a unique morphism

$$\Gamma_z : S_z \rightarrow \mathcal{G}_2 \otimes_{\mathcal{O}_2}^{\text{gp}} \Theta_z^{\text{st}}$$

such that:

- (i) *the derivative of Γ_z induces the natural inclusion of Ω_z^{st} into $\Omega_{S/A}$;*
- (ii) *the action of Φ on S_z is compatible with multiplication by p on $\mathcal{G}_2 \otimes_{\mathcal{O}_2}^{\text{gp}} \Theta_z^{\text{st}}$.*

Here, the variables on S_z obtained by pull-back via Γ_z are called the strong variables on S_z . Thus, these strong variables carry a Lubin-Tate geometry of order 2. Finally, the fiber of Γ_z over the identity element of the group object $\mathcal{G}_2 \otimes_{\mathcal{O}_2}^{\text{gp}} \Theta_z^{\text{st}}$ admits an anabelian geometry of length 2 and order 1 determined by Φ (plus a “Hodge subspace” for Φ – cf. Chapter IX, §1.5, for more details). The variables in these fibers are called the weak variables.

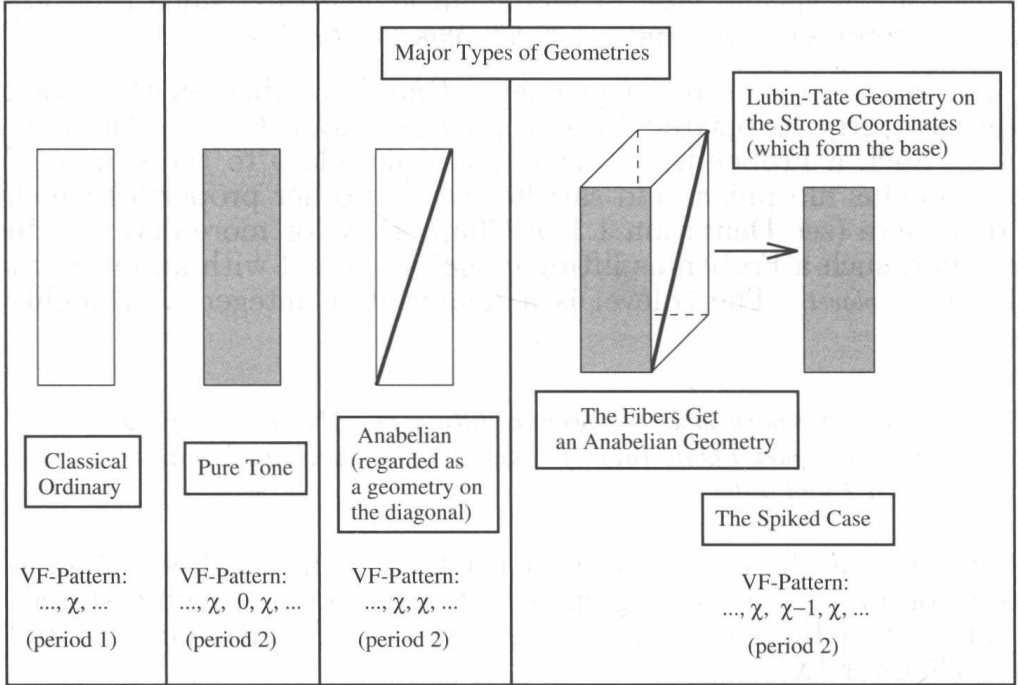


Fig. 15: Major Types of p -adic Geometries

Thus, in summary, Φ defines a *virtual fibration* on S to a base space (of dimension c) naturally equipped with a Lubin-Tate geometry of order 2; moreover, (roughly speaking) the fibers of this fibration are naturally equipped with an anabelian geometry of length 2 and order 1. In terms of VF-patterns, this sort of Frobenius lifting occurs in the case $\varpi = 2$, $\Pi(1) \neq 0$ (cf. Theorem 1.15). The colevel is then given by $2(\chi - \Pi(1))$.

Next, we note that as remarked toward the end of §1.5, in the binary ordinary and very ordinary spiked cases one obtains geometrizable systems of Frobenius liftings not only over $\mathcal{N}_{\mathbf{Z}_p}^{\text{ord}}$ (which is étale over $(\overline{\mathcal{M}}_{g,r})_{\mathbf{Z}_p}$) but also on the ordinary locus of the universal curve over $\mathcal{N}_{\mathbf{Z}_p}^{\text{ord}}$. (More precisely, in the very ordinary spiked case, one must replace $\mathcal{N}_{\mathbf{Z}_p}^{\text{ord}}$ by the formal open substack defined by the very ordinary locus.) Thus, in particular,

In the binary ordinary and very ordinary spiked cases, one obtains geometries as discussed in the above theorems not only on the moduli of the curves in question, but also on the ordinary loci of the universal curves themselves.

See Fig. 15 for a pictorial representation of the major types of geometries discussed.

Finally, we observe that one way to understand these generalized ordinary geometries is the following:

The “Lubin-Tate-ness” of the resulting geometry on the moduli stack is a reflection of the extent to which the p -curvature (of the indigenous bundles that the moduli stack parametrizes) vanishes.

That is to say, the more the p -curvature vanishes, the more Lubin-Tate the resulting geometry becomes. For instance, in the case of a Lubin-Tate geometry, the order of the Lubin-Tate geometry (cf. Theorem 1.16) corresponds precisely to the number of dormant crystalline bundles in a period (minus one). In the case of a spiked geometry, the number of “Lubin-Tate dimensions” is measured by the colevel. Moreover, this colevel is proportional to the degree of vanishing of the p -curvature of the indigenous bundle in question.

§1.7. The Canonical Galois Representation

Finally, since we have been considering Frobenius invariant indigenous bundles,

We would like to construct representations of the fundamental group of the curve in question into PGL_2 by looking at the Frobenius invariant sections of these indigenous bundles.

Such representations will then be the p -adic analogue of the canonical representation in the complex case of the topological fundamental group of a hyperbolic Riemann surface into $\mathrm{PSL}_2(\mathbf{R}) \subseteq \mathrm{PGL}_2(\mathbf{C})$ (cf. the discussion at the beginning of §0.3). Unfortunately, things are not so easy in the p -adic (generalized ordinary) case because *a priori* the canonical indigenous bundles constructed in Theorem 1.13 only have connections and Frobenius actions with respect to the *relative* coordinates of the tautological curve over $\mathcal{N}_{\mathbf{Z}_p}^{\mathrm{ord}}$. This means, in particular, that we cannot immediately apply the theory of [Falt1], §2, to pass to representations of the fundamental group. To overcome this difficulty, we must employ the technique of *crystalline induction* developed in [Mzk1]. Unfortunately, in order to carry out crystalline induction, one needs to introduce an object called the *Galois mantle* which can

only be constructed when the system of Frobenius liftings on $\mathcal{N}_{\mathbf{Z}_p}^{\text{ord}}$ is geometrizable. Thus, in particular, we succeed (in Chapter X) in constructing representations of the sort desired only in the binary ordinary and very ordinary spiked cases.

First, we sketch what we mean by the *Galois mantle*. The Galois mantle can be constructed for any geometrizable system of Frobenius liftings (e.g., any of the types discussed in §1.6). In particular, the notion of the Galois mantle has nothing to do with curves or their moduli. For simplicity, we describe the Galois mantle in the *classical ordinary case*. Thus, let S and A be as in §1.6. Let Π_S be the fundamental group of $S \otimes_{\mathbf{Z}_p} \mathbf{Q}_p$ (for some choice of base-point). Let Φ be a *classical ordinary Frobenius lifting* (in other words, Lubin-Tate of order 1) on S . Then by taking Frobenius invariant sections of the tangent bundle, one obtains an étale local system $\Theta_\Phi^{\text{ét}}$ on S of free \mathbf{Z}_p -modules of rank $\dim_A(S)$. Moreover, Φ defines a natural exact sequence of continuous Π_S -modules

$$0 \rightarrow \Theta_\Phi^{\text{ét}}(1) \rightarrow E_\Phi \rightarrow \mathbf{Z}_p \rightarrow 0$$

where the “(1)” denotes a Tate twist, and “ \mathbf{Z}_p ” is equipped with the trivial Π_S -action. Roughly speaking, this extension of Π_S -modules is given by taking the p^{th} power roots of the canonical multiplicative coordinates of Theorem 0.3 (cf. §2.2 of Chapter VIII for a detailed discussion of the p -divisible group whose Tate module may be identified with E_Φ). Let \mathcal{B}' be the affine space of dimension $\dim_A(S)$ over \mathbf{Z}_p parametrizing splittings of the above exact sequence. Then the action of Π_S on the above exact sequence induces a natural action of Π_S on \mathcal{B}' . Roughly speaking, the *Galois mantle* \mathcal{B} associated to Φ is the p -adic completion of a certain kind of p -adic localization of \mathcal{B} .

More generally, to any geometrizable system of Frobenius liftings (as in §1.6) on S , one can associate a natural p -adic space \mathcal{B} – the Galois mantle associated to the system of Frobenius liftings – with a continuous Π_S -action. In the binary ordinary case, \mathcal{B} will have a natural affine structure over some finite étale extension of \mathbf{Z}_p . In the very ordinary spiked case, \mathcal{B} will be fibred over an affine space over \mathcal{O}_2 with fibers that are also equipped with an affine structure over \mathcal{O}_2 .

In fact, to be more precise, \mathcal{B} is only equipped with an action by a certain open subgroup of Π_S , but we shall ignore this issue here since it is rather technical and not so important. We refer to §2.3 and §2.5 of Chapter IX for more details on the Galois mantle. So far, for simplicity, we have been ignoring the logarithmic case, but everything is compatible with log structures.

We are now ready to state the main result on the canonical Galois representation in the generalized ordinary case, i.e., the generalized ordinary analogue of Theorem 0.4 (cf. Theorems 1.2 and 2.2 of Chapter X). See Fig. 16 for a graphic depiction of this theorem.




Fig. 16: The Canonical Galois Representation

Theorem 1.19. (Canonical Galois Representation) *Let p be an odd prime. Let g and r be nonnegative integers such that $2g - 2 + r \geq 1$. Fix a VF-pattern (Π, ϖ) which is either binary ordinary or spiked of order 2. Let $S \stackrel{\text{def}}{=} \mathcal{N}_{\mathbf{Z}_p}^{\text{ord}}$ in the binary ordinary case, and let S be the very ordinary locus of $\mathcal{N}_{\mathbf{Z}_p}^{\text{ord}}$ in the spiked case. Let $Z \rightarrow S$ be a certain appropriate finite covering which is log étale in characteristic zero (cf. the discussion preceding Theorems 1.2 and 2.2 of Chapter X for more details). Let $X_Z^{\log} \rightarrow Z^{\log}$ be the tautological log-curve over Z^{\log} . Let Π_{X_Z} (respectively, Π_Z) be the fundamental group of $X_Z^{\log} \otimes_{\mathbf{Z}_p} \mathbf{Q}_p$ (respectively, $Z^{\log} \otimes_{\mathbf{Z}_p} \mathbf{Q}_p$) for some choice of base-point. (Of course, despite the similarity in notation, these fundamental groups have no direct relation to the VF-pattern “ Π .”) Thus, there is a natural morphism $\Pi_{X_Z} \rightarrow \Pi_Z$. Let \mathcal{B} be the Galois mantle associated to the canonical system of Frobenius liftings of Theorem 1.13. The morphism $\Pi_{X_Z} \rightarrow \Pi_Z$ allows us to regard \mathcal{B} as being equipped with a Π_{X_Z} -action.*

Let \mathcal{P} be the tautological Π -indigenous bundle on X . Then by taking Frobenius invariants of \mathcal{P} , one obtains a \mathbf{P}^1 -bundle

$$\mathbf{P}_{\mathcal{B}} \rightarrow \mathcal{B}$$

equipped with a natural continuous Π_{X_Z} -action compatible with the above-mentioned action of Π_{X_Z} on the Galois mantle \mathcal{B} .

Put another way, one obtains a twisted homomorphism of Π_{X_Z} into PGL_2 of the functions on \mathcal{B} . (Here, “twisted” refers to the fact that the multiplication rule obeyed by the homomorphism takes into account the action of Π_{X_Z} on the functions on \mathcal{B} .) Finally, note that for any point of $Z \otimes_{\mathbf{Z}_p} \mathbf{Q}_p$ (at which the log structure is trivial), one also obtains similar representations by restriction. This gives one canonical Galois representations even in the non-universal case. Finally, in Chapter X, §1.4, 2.3, we show that:

The Galois representation of Theorem 1.19 allows one to relate the various p -adic analytic structures constructed throughout the work (i.e., canonical Frobenius liftings, canonical Frobenius invariant indigenous bundles, etc.) to the algebraic/arithmetical Galois action on the profinite Teichmüller group (cf. Chapter X, Theorems 1.4, 2.3).

This results in a rather detailed understanding of a certain portion of a canonical tower of coverings of $(\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{Z}_p}$, analogous to the analysis given in [KM] of coverings of the moduli stack of elliptic curves over \mathbf{Z}_p obtained by considering p -power torsion points (cf. the Remark following Chapter X, Theorem 1.4, for more details). The text ends with a concluding section (Chapter X, §3) that ties together the Galois representations of Chapter X with the corresponding objects in the complex theory, as well as with much of the theory of the rest of the text.

Thus, in summary, Theorem 1.19 concludes our discussion of “ p -adic Teichmüller theory” as exposed in this book by constructing a p -adic analogue of the canonical representation discussed at the beginning of §0.3, that is to say, a p -adic analogue of something very close to the Fuchsian uniformization itself – which was where our discussion began (§0.1).

§1.8. Ordinary Stable Bundles

Finally, we included an Appendix concerning “ordinary” stable bundles on a curve. This theory of ordinary stable bundles is the theory whose existence is referred to in the discussion of §0.8. In particular, we discuss two types of uniformization result:

- (1) we show that “any” stable bundle can be constructed from a *representation of the fundamental group* of the curve into some compact subgroup of GL_r (Theorems 2.1 and 3.12 of the Appendix);
- (2) we show the existence of *canonical affine uniformizations* of the moduli space of stable bundles that generalize the exponential map in the abelian case (Theorems 2.5 and 3.16 of the Appendix).

Moreover, it is observed that both types of uniformization exist in both the complex and p -adic cases and that the complex and p -adic theories are substantially analogous to one another.

Unlike the rest of this book, this Appendix is far from original and merely reformulates many well-known classical results in a form compatible with the philosophy of the present book. Moreover, this theory has nothing *logically* to do with the “ p -adic Teichmüller theory” of [Mzk1] or the present book. However, (as discussed in §0.8) it gives an interesting look at one more case of the asserted analogy between Frobenius actions and Kähler metrics. Compared to the extensive theory of [Mzk1] and the present book on the moduli of curves, this theory of stable bundles is technically much simpler, hence much easier to grasp for the first-time reader. Moreover, historically, the author first discovered (without bothering to publish – since it seemed to have little “original” content) this theory of ordinary stable bundles in the Summer of 1993, just before he started to work on the theory of [Mzk1]. That is to say,

The theory of ordinary stable bundles formed the prototype for the theory of [Mzk1], as well as the theory of the present book, and hence may provide the reader with useful psychological insights concerning what the author has been trying to achieve in [Mzk1] and the present book.

§2. Open Problems

So far, we have been discussing what is *known* concerning the theory discussed in this book. In the present, final § of this introductory chapter, we would like to discuss what the author believes to be the main open questions in this theory. Also, we feel that seeing what is still unknown may help the reader to get a better sense of the boundaries and limitations of the (known) theory discussed in this book.

§2.1. Basic Questions

In some sense, the most basic object in the theory of [Mzk1] and the present book is the *moduli stack of nilcurves* $\overline{\mathcal{N}}_{g,r}$ equipped with its natural morphism

$$\overline{\mathcal{N}}_{g,r} \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$$

to the moduli stack of pointed stable curves in characteristic p . As discussed earlier (Theorem 0.1), this morphism is a finite, flat, local complete intersection morphism of degree p^{3g-3+r} . We denote the étale locus of this morphism by

$$(\overline{\mathcal{N}}_{g,r}^{\text{ord}})_{\mathbf{F}_p}$$

Then one can ask the following natural questions:

- (1) *Is the natural morphism $(\overline{\mathcal{N}}_{g,r}^{\text{ord}})_{\mathbf{F}_p} \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ surjective? Of course, the analogy with ordinary abelian varieties might lead one to suspect that the answer to this question is negative, but already in the case of $g = 1$, $r = 1$, $p = 5$ (cf. Chapter IV, §1.3), one can in fact prove surjectivity.*
- (2) *Suppose that X is a pointed stable curve equipped with a nilpotent ordinary indigenous bundle \mathcal{P} . Let $Y \rightarrow X$ be a finite (log) étale covering of X . Then is the pull-back of \mathcal{P} to Y still ordinary? This question appears to be related to the issue of the closedness of the ordinary locus in \mathcal{Q}^Π , i.e.: For, say, a binary VF-pattern Π , is \mathcal{Q}^{ord} closed (i.e., not just ω -closed – cf. Theorem 1.12) in \mathcal{Q}^Π ? In an earlier (preprint) version of this work (cf. [Mzk5,6]), it was stated that the binary Π -ordinary locus is closed (i.e., not just ω -closed), and this closedness result was applied to prove a result concerning the extent to which ordinarity was preserved under passing to finite log étale coverings. Upon closer inspection, however, the author found a gap in the proof of this closedness result – a gap which showed that in fact, what he had really proved was just “ ω -closedness” not “closedness.” Thus, the question of closedness remains open. Moreover, in the case of more general (i.e., non-binary) VF-patterns, at the time of writing, one does not even have a proof of the ω -closedness of the Π -ordinary locus.*
- (3) *As discussed in Theorem 0.3, there is a canonical local system Ω_Φ^{et} on $(\overline{\mathcal{N}}_{g,r}^{\text{ord}})_{\mathbf{F}_p}$ of free \mathbf{Z}_p -modules of rank $3g - 3 + r$. What does*

the monodromy group of this local system look like? Moreover, for each individual k -valued point of $(\overline{\mathcal{N}}_{g,r}^{\text{ord}})_{\mathbb{F}_p}$ (where k is a finite field), one has an action of Frobenius on the \mathbb{Z}_p -module $\Omega_{\Phi}^{\text{et}}|_{\overline{k}}$ (where \overline{k} is an algebraic closure of k). What can one say about the eigenvalues of this Frobenius endomorphism? Are these eigenvalues algebraic numbers (as in the case of the Weil Conjectures)? At the present time, I have no idea what the answer to these questions is, but I think that these questions are important and fundamental. Of course, although in this paragraph, we have only been discussing the classical ordinary case, there is also a generalized ordinary analogue of this monodromy question.

- (4) What can one say about the irreducible components of $\mathcal{N}_{g,r}$? Can one give a modular interpretation to the decomposition into irreducible components? At the present time, one has a combinatorial algorithm for computing the degree of each $\mathcal{N}_{g,r}[d]$ over $(\mathcal{M}_{g,r})_{\mathbb{F}_p}$ (cf. Theorem 1.4), but one has no idea what the irreducible components of each $\mathcal{N}_{g,r}[d]$ look like. One might naively conjecture that each $\mathcal{N}_{g,r}[d]$ is already irreducible, but explicit computations in the case $g = 1$, $r = 1$, $p = 5$ show that this is false (cf. Chapter IV, §1.3). In some sense, this question concerning irreducible components is also a sort of “monodromy question,” hence perhaps it should be studied in the context of Question (3) above.
- (5) Can one come to a more meaningful understanding of the complicated combinatorics associated with Theorem 1.4? For instance, how can one interpret the strictly combinatorial identities that arise as consequences of Theorem 1.4? Also, are the $n_{g,r,p}^{\text{ord}}$ ’s really polynomials in p , and if so, can one give a meaningful recipe for computing the coefficients of these polynomials? Although at the present time one lacks even a conjectural answer to these questions, various explicit computations in special cases are given in Chapter V, §3.2.
- (6) Can one use the theory of nilcurves to construct natural proper subvarieties of $\mathcal{N}_{g,r}$? Can one derive from the existence of such proper subvarieties consequences concerning the minimal codimension of proper subvarieties of $\mathcal{M}_{g,r}$? This question was proposed to the author by Prof. Frans Oort. It is motivated by the situation in the case of abelian varieties, where there do exist natural proper subvarieties (of the moduli stack of principally polarized abelian varieties) which allow one to prove results concerning the minimal codimension of proper subvarieties of this stack (cf. Corollary 1.7 of [Geer]; [Oort]).

- (7) *Theorem 0.3 gives rise to a theory of canonical liftings of ordinary nilcurves over perfect fields (cf. Chapter IV of [Mzk1]). What interesting properties do such canonically lifted curves possess? Are such curves defined over a number field? Of course, there is also a generalized ordinary analogue of this question. At the present time, the author has no idea what the answer to these questions is, but the results of [Mzk4] lead one to be rather pessimistic about the possibility that such canonical curves are always defined over a number field. We will take a closer look at this question in the following §.*
- (8) *In §1.6, we discussed how to “geometrize” systems of Frobenius liftings for certain special types of VF-pattern Π . Can one generalize this to obtain a theory of geometrization of the sorts of systems of Frobenius liftings that arise in Theorem 1.13 for arbitrary VF-patterns Π ? This question appears to be related to the issue of showing ω -closedness for more general VF-patterns Π (cf. Question (2) above).*

§2.2. Canonical Curves and Hyperbolic Geometry

One approach to trying to understand canonical curves is to ask what the analogue of such curves is in the complex case. Indeed, in the case of elliptic curves (i.e., the case $g = 1, r = 0$), the key observation is that *Serre-Tate canonical curves admit complex multiplication*. This observation allows one to prove many interesting things concerning such curves (for instance, that they are defined over a number field). Also, over the complex numbers, it is very easy to understand what it means for an elliptic curve to admit complex multiplication. Thus, it is natural to ask whether an analogous theory exists for hyperbolic curves over the complex numbers, i.e.,

Is there a natural notion of “canonical” hyperbolic curves over the complex numbers?

One naive answer to this question is to consider curves that admit a correspondence, which is a sort of “hyperbolic isogeny.” The problem with this proposal is that (by a result of [Take], communicated to the author by Prof. Y. Ihara), there exist only finitely many hyperbolic curves (for a given g and r) that admit correspondences (cf. [Mzk4]). Since, on the other hand, there are countably infinitely many canonical curves, one thus sees that this proposal for an analogy cannot possibly be correct. Thus, in this §, we give a new proposal for an analogy based on three-dimensional real hyperbolic geometry.

§2.2.1. Review of Kleinian Groups

Let us recall the general set-up for Kleinian groups Γ (see [Mask], [Thur] for more details). One starts with a discrete subgroup $\Gamma \subseteq PGL_2(\mathbf{C})$, which defines a domain of discontinuity $D_\Gamma \subseteq \mathbf{P}_\mathbf{C}^1$. On D_Γ , Γ acts discontinuously, so we may form the quotient $R_\Gamma \stackrel{\text{def}}{=} D_\Gamma/\Gamma$. If Γ is *finitely generated*, then, by *Ahlfors' Finiteness Theorem*, R_Γ is a finite union of Riemann surfaces of finite type (which will typically be hyperbolic). In particular, for each connected Riemann surface of finite type $X \subseteq R_\Gamma$, we thus obtain an embedding $\tilde{X} \hookrightarrow \mathbf{P}_\mathbf{C}^1$ (where $\tilde{X} \rightarrow X$ is the universal covering). Just as in the construction of the canonical indigenous bundle (cf. the Introduction of [Mzk1]), one can thus construct an indigenous bundle on X from this embedding $\tilde{X} \hookrightarrow \mathbf{P}_\mathbf{C}^1$.

The Kleinian group Γ is called *quasi-Fuchsian* if D_Γ has exactly two connected components. Quasi-Fuchsian groups are integrally related to *simultaneous uniformizations of Riemann surfaces*, as follows: Let $f : X \rightarrow X'$ be a quasiconformal homeomorphism between two compact hyperbolic Riemann surfaces. Then, by passing to universal covering spaces, we obtain a quasiconformal map $\tilde{f} : \tilde{X} \rightarrow \tilde{X}'$. By choosing holomorphic isomorphisms $L \cong \tilde{X}$ and $L \cong \tilde{X}'$ (where $L \subseteq \mathbf{C}$ is the upper half-plane of \mathbf{C}), we may think of \tilde{f} as a map from L to L . Let μ be the *dilatation* (for a definition of this, we refer to [Lehto]) of \tilde{f} . Thus, \tilde{f} forms a Beltrami differential on L . Then it follows from the theory of quasiconformal maps that there exists a unique quasiconformal homeomorphism $\Phi : \mathbf{P}_\mathbf{C}^1 \rightarrow \mathbf{P}_\mathbf{C}^1$ whose dilatation is equal to zero on the upper half-plane H (i.e., it is holomorphic on H), and μ on L . Note that the canonical representation $\rho : \pi_1(X) \rightarrow PSL_2(\mathbf{R})$ defines a natural action of $\pi_1(X)$ on $\mathbf{P}_\mathbf{C}^1$. Moreover, it is not difficult to show that every element of the group of homeomorphisms Γ on $\mathbf{P}_\mathbf{C}^1$ obtained by conjugating (the action on $\mathbf{P}_\mathbf{C}^1$ of) $\pi_1(X)$ by Φ is holomorphic. Thus, $\Gamma \subseteq PGL_2(\mathbf{C})$, and Γ is a *quasi-Fuchsian* (hence, in particular, a Kleinian) group. Indeed, $D_\Gamma = \Phi(H) \cup \Phi(L)$. Finally, it is known that every quasi-Fuchsian group is obtained in this fashion.

§2.2.2. Review of Three-Dimensional Hyperbolic Geometry

Here, we follow the treatment of [Mask]. Let

$$\mathbf{H}^3 \stackrel{\text{def}}{=} \{(x, t) \in \mathbf{R}^2 \times \mathbf{R} | t > 0\}$$

We shall write \mathbf{E}^2 for the subset of $\mathbf{R}^2 \times \mathbf{R}$ defined by $t = 0$. We shall write $\partial\mathbf{H}^3$ for $\mathbf{E}^2 \cup \infty$. Often, we shall regard $\partial\mathbf{H}^3$ as equal to $\mathbf{P}_\mathbf{C}^1$. \mathbf{H}^3 is equipped with a natural *Riemannian metric*:

$$ds^2 = \frac{1}{t^2}(dx^2 + dt^2)$$

Next, we would like to define a *group of motions* acting on \mathbf{H}^3 (and $\partial\mathbf{H}^3$). There are four basic types of motions: translations (of the x -coordinate), rotations (of the x -coordinate), dilations (of both x and t), and inversions. (We refer to [Mask] for a detailed treatment of these motions). One checks easily that these motions all preserve the metric, and, in fact, generate the group of isometries of (\mathbf{H}^3, ds^2) . This group of isometries can be shown to be naturally isomorphic to $PGL_2(\mathbf{C})$. Moreover, the induced action of $PGL_2(\mathbf{C})$ on the boundary $\partial\mathbf{H}^3 = \mathbf{P}_{\mathbf{C}}^1$ is the usual action of $PGL_2(\mathbf{C})$ on $\mathbf{P}_{\mathbf{C}}^1$.

The action of $PGL_2(\mathbf{C})$ on \mathbf{H}^3 is transitive on points and tangent vectors (i.e., transitive on the total space of the sphere bundle associated to the tangent bundle of \mathbf{H}^3). Note that one has a copy of $SO(3) = SU(2)/\{\pm 1\} \subseteq PGL_2(\mathbf{C})$. It is not difficult to show that $SO(3)$ is a maximal compact subgroup of $PGL_2(\mathbf{C})$, and that, moreover, the isotropy subgroup of the action of $PGL_2(\mathbf{C})$ on \mathbf{H}^3 at any point of \mathbf{H}^3 is conjugate to $SO(3)$. Thus,

$$\mathbf{H}^3 = PGL_2(\mathbf{C})/SO(3)$$

The geodesics of \mathbf{H}^3 are arcs of circles (and lines) orthogonal to \mathbf{E}^2 . A typical geodesic is given by the line $(0, 0, t)$. The isotropy subgroup of a geodesic is naturally isomorphic to \mathbf{C}^\times .

A *complete (real) three-dimensional hyperbolic manifold* is defined to be a Riemannian manifold which is isomorphic to \mathbf{H}^3/Γ , where $\Gamma \subseteq PGL_2(\mathbf{C})$ is a discrete subgroup with no fixed points. (Of course, one can also consider the case with fixed points, by working with stacks.) An alternate characterization of such manifolds is that they are complete Riemannian manifolds whose curvature is identically equal to -1 .

Let $\Gamma \subseteq PGL_2(\mathbf{C})$ be a Kleinian group. Then, on the one hand, in the preceding subsection, we formed a Riemann surface $R_\Gamma = D_\Gamma/\Gamma$, where $D_\Gamma \subseteq \mathbf{P}_{\mathbf{C}}^1$. On the other hand, we can also form the quotient $M_\Gamma \stackrel{\text{def}}{=} \mathbf{H}^3/\Gamma$, which is called a *Kleinian three-fold*. Thus, M_Γ is a complete hyperbolic manifold. Moreover, since $\mathbf{P}_{\mathbf{C}}^1$ may be regarded as the boundary of \mathbf{H}^3 , by passing to quotients, one sees that R_Γ may be regarded as the boundary of the open manifold M_Γ . By gluing on this boundary, we thus obtain a manifold with boundary \overline{M}_Γ .

In general, Kleinian groups have nontrivial moduli. Thus, if we deform Γ , we obtain deformations of R_Γ and M_Γ . A theorem of Ahlfors, Bers, Sullivan, and others states that if Γ is *finitely generated*, then relative to this correspondence, the so-called “quasi-isometric” deformations of the hyperbolic manifold M_Γ are in one-to-one correspondence

with the deformations of the conformal structure on R_Γ . Thus, in particular, one sees that Kleinian three-folds tend to have nontrivial continuous moduli.

§2.2.3. Rigidity and Density Results

Let M be a complete hyperbolic three-fold. Since M is equipped with a metric, it is also equipped, in particular, with a volume form, which we can integrate over M to form $v(M)$, the *volume* of M . In general, $v(M)$ can be finite or infinite. In the finite case, one has the following result:

Theorem 2.1. (Rigidity Theorem of Mostow-Prasad) *Let M_1 and M_2 be complete hyperbolic three-folds of finite volume. Let us consider the following sorts of objects: (i) isometries $f : M_1 \rightarrow M_2$; (ii) homotopy equivalences $f : M_1 \rightarrow M_2$; (iii) group isomorphisms $\pi_1(M_1) \rightarrow \pi_1(M_2)$. Then the natural maps that assign objects of type (ii) to objects of type (i), and objects of type (iii) to objects of type (ii) are bijections.*

(In fact, this holds in dimensions ≥ 3 , but we are only interested in the three-dimensional case.) Note that the behavior of complete hyperbolic three-folds described in this Theorem differs substantially from the two-dimensional case. Indeed, (as noted in the introduction to [Brooks]) one consequence of the above Rigidity Theorem is the following:

Corollary 2.2. (Countability) *There are only countably many hyperbolic three-folds of finite volume.*

This follows from Theorem 2.1 because by Theorem 2.1, one sees that it suffices to show that only countably many homotopy equivalence classes of M 's can occur. But this is not difficult to prove.

On the other hand, in addition to the above countability result, one also has the following density result: If Γ is a Kleinian group, and there exists a Kleinian group Γ^* such that $\Gamma \subseteq \Gamma^*$, and \mathbb{H}^3/Γ^* has finite volume, then we shall say (following [Brooks]) that Γ *admits a cofinite extension*. Then Brooks has proven the following result (Theorem 2 of [Brooks]):

Theorem 2.3. (Existence of Deformations with Cofinite Extension) *Let Γ be a geometrically finite Kleinian group. Then there exist arbitrarily small quasi-conformal deformations Γ_ϵ of Γ , such that Γ_ϵ admits a cofinite extension.*

We will not discuss the proof of Theorem 2.3 here, but the proof (which involves considering circle packing problems on R_Γ) is interesting in and of itself.

§2.2.4. QF-Canonical Curves

Now let us call a hyperbolic Riemann surface X (of finite type) *QF-canonical* (here “QF” stands for “quasi-Fuchsian”) if X appears as a connected component of R_Γ for a quasi-Fuchsian group Γ that admits a cofinite extension. Thus, Corollary 2.2 and Theorem 2.3 imply that:

Corollary 2.4. (Countability and Density) *The QF-canonical hyperbolic Riemann surfaces of finite type are countably infinite in number, but dense in $(\mathcal{M}_{g,r})_{\mathbb{C}}$.*

It is these QF-canonical curves that we wish to propose as possible analogues at the infinite prime of the p -adic canonical curves of [Mzk1] and the present book. Indeed, both are countable in number, but dense (in some appropriate topology) in $\mathcal{M}_{g,r}$.

We remark that there are various alternative definitions that one could take for “canonical curves at the infinite prime.” For instance, one could replace “quasi-Fuchsian” in the definition of “QF-canonical” by “Kleinian.” In fact, there are many other ways to modify the definition slightly. We chose the definition that we did not because we have any reason to believe strongly that it is the final answer to the question of finding analogues at the infinite prime of p -adic canonical curves, but because we thought that it was representative of the various possible candidate definitions that one could give.

Finally, let us note the following. Although, as discussed earlier, the hyperbolic curves which admit (algebraic) correspondences are only finite in number, hence cannot comprise the entire set of canonical curves at the infinite prime, it is, nonetheless, natural to require that hyperbolic curves which admit correspondences (like Shimura curves) be canonical at the infinite prime, since they are canonical at finite primes (as long as they are ordinary). This appears generally to be the case (with respect to the above definition of QF-canonical). Indeed, for instance, in the case of $\mathcal{M}_{1,1}$, which, over \mathbb{C} , is obtained from the Kleinian group $\Gamma = SL_2(\mathbb{Z})$, we have $\Gamma \subseteq \Gamma^* \stackrel{\text{def}}{=} SL_2(\mathbb{Z}[i])$, and, moreover, $SL_2(\mathbb{C})/\Gamma^*$ is of finite volume (by Theorem 7.8 of [BHC]) – a fact pointed out to the author by Prof. Y. Ihara – from which it follows immediately that $v(H/\Gamma^*) < \infty$.

§2.2.5. The Case of CM Elliptic Curves

To understand further why this sort of definition for a canonical curve at the infinite prime is natural, let us review the case of CM

elliptic curves. Let E be a compact Riemann surface of genus 1, with a given base point. Then we may write $E = \mathbf{C}/\Lambda$, where $\Lambda \cong \mathbf{Z}^2$ is a lattice in \mathbf{C} . For each one parameter subgroup $H \subseteq \mathbf{C}^\times$ such that $H \cong \mathbf{R}$ and $H \not\subseteq \mathbf{R}^\times$, let G_H be the three-dimensional (over \mathbf{R}) Lie group defined by taking the semi-direct product of H with \mathbf{C} , where H acts on \mathbf{C} by $H \hookrightarrow \mathbf{C}^\times$. Thus, we have an exact sequence of Lie groups

$$1 \rightarrow \mathbf{C} \rightarrow G_H \rightarrow H \rightarrow 1$$

Moreover, one sees easily that:

Proposition 2.5. (The Case of Elliptic Curves) *E admits CM if and only if there exists some H as above such that the above exact sequence surjects onto an exact sequence*

$$1 \rightarrow E_{\mathbf{Q}} \stackrel{\text{def}}{=} \mathbf{C}/\Lambda_{\mathbf{Q}} \rightarrow M \stackrel{\text{def}}{=} G_H/\Xi \rightarrow H/P \rightarrow 1$$

where $\Lambda_{\mathbf{Q}} = \Lambda \otimes_{\mathbf{Z}} \mathbf{Q} \cong \mathbf{Q}^2$; $P \cong \mathbf{Z}$ (so $H/P \cong S^1$); and Ξ is an extension of P by $\Lambda_{\mathbf{Q}}$.

In other words, this situation is analogous to the definition of QF-canonical curves in the hyperbolic case in that the condition on E is that there exist an extension $\Lambda_{\mathbf{Q}} \subseteq \Xi$ which is “cofinite” in the sense that G_H/Ξ is essentially “compact” (i.e., relative to regarding “ $E_{\mathbf{Q}}$ ” as “compact”), i.e., of finite volume.

Thus, Proposition 2.5 is a *geometric* characterization of the property of having CM. Note that the “geometry” here is the three-dimensional geometry based on the Lie group G_H . It is not difficult to compute the geodesics in this geometry. We leave this task to the reader.

One important (possible) difference between the parabolic and hyperbolic cases, however, is the following: In the parabolic case, it is immediate that elliptic curves that satisfy the condition of Proposition 2.5 at the infinite prime, as well as Serre-Tate-canonical elliptic curves are *defined over an algebraic number field*. On the other hand, it seems to the author that there is no reason to believe that either QF-canonical hyperbolic Riemann surfaces or p -adic canonical hyperbolic curves (as in [Mzk1] and the present book) should be defined (in general) over an algebraic number field. (At the present time, however, the author knows of no counterexamples.)

§2.2.6. The Third Real Dimension as the Frobenius Dimension

So far, (to the author’s knowledge), there have not been any connections established between the geometry of (real) *three-dimensional*

manifolds and algebraic geometry. Indeed, the odd number of real dimensions seems to fly in the face of any attempt to relate such manifolds to the world of algebraic geometry, where it seems to be of absolute necessity that dimensions come in two's. Thus, in this subsection, we would like to discuss why the relationship proposed above between canonical hyperbolic curves and certain hyperbolic three-folds is, in fact, natural.

The point is that the third “nonalgebraic” dimension should be regarded as “ $\mathbf{R} \cdot Fr$,” i.e., as a one-dimensional real vector space generated by the symbol “Frobenius.” Alternatively, one may regard the third dimension as being the value group of the valued field \mathbf{C} . From the point of view of class field theory, this essentially amounts to the same thing, since, in class field theory, at finite primes, the value group is mapped isomorphically to the free subgroup of the Galois group of the residue field generated by Frobenius.

Thus, the point is that canonical curves X (at both finite and infinite primes) are essentially curves that occur in objects Z which are twisted products of the curve plus a copy of $\mathbf{R} \cdot Fr$ or $\mathbf{Z}_p \cdot Fr$.

In the complex case, Z is \mathbf{H}^3/Γ^* (where Γ^* is the cofinite extension of the quasi-Fuchsian group Γ for which $X \subseteq R_\Gamma$), and the extra Frobenius dimension is the dimension corresponding to the coordinate “ t ” in the definition of \mathbf{H}^3 . Thus, in particular, the dimension of Z over \mathbf{R} works out to be 3. In the p -adic case, a “ Z ” in some sort of formal, hypothetical category can be obtained by taking the twisted product of $\mathbf{Z}_p \cdot Fr$ with X , where one lets Fr “act” on X via the “quasi-analytic self-isogeny” on X given by the canonical indigenous bundle on X .

Thus, relative to this point of view,

There is a certain analogy between Thurston’s program (see, e.g., [Thurs]) of geometrizing all real three-dimensional manifolds and the aim of the present work, namely, to classify (and geometrize! – cf. §1.6) all Frobenius invariant indigenous bundles (i.e., “three-dimensional p -adic objects” obtained by taking the twisted product of a hyperbolic curve with $\mathbf{Z}_p \cdot Fr$, where the twisting is defined by the “quasi-analytic self-isogeny” corresponding to the Frobenius invariant indigenous bundle).

Indeed, an important example of the geometrization of three-dimensional manifolds that occurs in Thurston’s program is precisely the example of a twisted product of S^1 (i.e., the circle) with a compact hyperbolic surface, where the twist is defined by a quasi-conformal automorphism of the surface. Another interesting piece of circumstantial evidence for an analogy between the p -adic and real theories is the formal similarity between the Rigidity Theorem (Theorem 2.1) and the Grothendieck Conjecture (cf. §0.10, especially Theorem 0.5).

§2.3. Towards an Arithmetic Kodaira-Spencer Theory

In this §, we give another proposal for another sort of fundamental problem concerning the *global* arithmetic of curves. Here, by “global,” we mean that it relates to all the primes of a number field. The problem is the following: Suppose that S is a smooth, proper, connected curve over an algebraically closed field k . Suppose that $f : X \rightarrow S$ is a family of stable curves of genus ≥ 2 . Then we can form the Kodaira-Spencer morphism

$$\kappa : f_*\omega_{X/S}^{\otimes 2} \rightarrow \omega_{S/k}$$

of f , which basically amounts to taking the derivative of the classifying morphism $S \rightarrow \overline{\mathcal{M}}_g$ defined by f . The morphism κ measures (at least in characteristic 0) the extent to which $f : X \rightarrow S$ departs from being an isotrivial family (i.e., a family which becomes trivial after passing to a finite, flat covering of S), or, more concretely, from being *constant*. The existence of the Kodaira-Spencer morphism is very useful in studying the arithmetic of curves over S . For instance, as in [EV], the existence of the Kodaira-Spencer morphism is what allows one to prove “Vojta’s conjecture” (see, e.g., [Vojta]) in the geometric case.

On the other hand, despite the generally assumed analogy between number fields and function fields, there is no known analogue of the Kodaira-Spencer morphism κ for a curve of genus ≥ 2 over a number field. This is closely related to the problem that one is not sure how to define a “constant” curve over a number field F , i.e., a curve $X \rightarrow \operatorname{Spec}(F)$ that acts as if it is obtained by base-change from some field of constants inside F . Another closely related problem is that one doesn’t know how to compare curves over fields of different characteristics. We shall refer to this collection of issues as the *arithmetic Kodaira-Spencer problem*. Considering the situation in the geometric case (i.e., [EV]), if one could satisfactorily solve the arithmetic Kodaira-Spencer problem, then it is not unreasonable to hope that one might be able to prove some version of Vojta’s conjecture for curves over number fields. It is in this sense that the arithmetic Kodaira-Spencer problem is of interest not just theoretically or philosophically, but also from a “practical” point of view, as well.

In this §, we discuss how the theory of [Mzk1] and the present book appear to bear some relation to the arithmetic Kodaira-Spencer problem.

§2.3.1. The Schwarz Torsor as Dual to the Kodaira-Spencer Morphism

Since one doesn’t know how to “differentiate” a morphism $\operatorname{Spec}(F) \rightarrow \overline{\mathcal{M}}_{g,r}$ (where F is a number field), it is natural to try to translate

the Kodaira-Spencer morphism in the geometric case into some sort of more abstract form which will have a more obvious analogue in the number field case. To do this, let us work (for simplicity) over \mathbf{C} , and consider the *Schwarz torsor* (of [Mzk1], Chapter I)

$$\overline{\mathcal{S}}_{g,r} \rightarrow \overline{\mathcal{M}}_{g,r} = (\overline{\mathcal{M}}_{g,r})_{\mathbf{C}}$$

Since this is a torsor over $\Omega_{\overline{\mathcal{M}}_{g,r}/\mathbf{C}}^{\log}$, it defines a “Schwarz class”:

$$\Sigma_{g,r} \in H^1(\overline{\mathcal{M}}_{g,r}, \Omega_{\overline{\mathcal{M}}_{g,r}/\mathbf{C}}^{\log})$$

Recall that the Schwarz torsor is “compactly supported” over $\mathcal{M}_{g,r}$ in the sense that, “the more singular the curve, the more of a canonical trivialization (given by considering *indigenous bundles of restricted type*) one has.” Thus, by using these canonical trivializations, one finds that in fact, one has a compactly supported class

$$\Sigma_{g,r}^c \in H_c^1(\overline{\mathcal{M}}_{g,r}, \Omega_{\overline{\mathcal{M}}_{g,r}/\mathbf{C}}^{\log})$$

Here, $H_c^1(\overline{\mathcal{M}}_{g,r}, \Omega_{\overline{\mathcal{M}}_{g,r}/\mathbf{C}}^{\log})$ is defined as follows: Let $\mathcal{R} \stackrel{\text{def}}{=} \Omega_{\overline{\mathcal{M}}_{g,r}/\mathbf{C}}^{\log} / \Omega_{\overline{\mathcal{M}}_{g,r}/\mathbf{C}}$. Then $H_c^i(\overline{\mathcal{M}}_{g,r}, \Omega_{\overline{\mathcal{M}}_{g,r}/\mathbf{C}}^{\log})$ is defined as the hypercohomology of the complex (in degrees 0 and 1 where the morphism is the natural projection):

$$\Omega_{\overline{\mathcal{M}}_{g,r}/\mathbf{C}}^{\log} \rightarrow \mathcal{R}$$

That is, in this case, the compact cohomology $H_c^i(\overline{\mathcal{M}}_{g,r}, \Omega_{\overline{\mathcal{M}}_{g,r}/\mathbf{C}}^{\log})$ turns out to be simply $H^i(\overline{\mathcal{M}}_{g,r}, \Omega_{\overline{\mathcal{M}}_{g,r}/\mathbf{C}})$. It was shown in §3 of [Mzk1], Chapter I, that in general, $\Sigma_{g,r}$ is highly nontrivial, and hence, *a fortiori*, so is $\Sigma_{g,r}^c$.

Now let S be a smooth, proper, connected curve over \mathbf{C} , and suppose that we are given a morphism $\phi : S \rightarrow \overline{\mathcal{M}}_{g,r}$ whose image has nonempty intersection with $\mathcal{M}_{g,r}$. Let $B \subseteq S$ be the closed reduced subscheme consisting of those points that are sent by ϕ to points of $\overline{\mathcal{M}}_{g,r} - \mathcal{M}_{g,r}$. Thus, B is a divisor with normal crossings on S , so it defines a log scheme S^{\log} . Moreover, we have a residue map $\text{Res} : \omega_{S^{\log}/\mathbf{C}} \rightarrow \mathcal{O}_B$ given by taking the residues of differentials at the points of B . We shall write $H_c^1(S, \omega_{S^{\log}/\mathbf{C}})$ for the hypercohomology of the complex defined by Res . Note that $H_c^1(S, \omega_{S^{\log}/\mathbf{C}}) = H^1(S, \omega_S/\mathbf{C}) = \mathbf{C}$.

We would like to pull the class $\Sigma_{g,r}^c$ back to S via Φ . Since we want to continue to work with compact supports, we consider the complex

$$\Omega_{\mathcal{M}_{g,r}/\mathbf{C}}^{\log}|_S \rightarrow \mathcal{R}|_B$$

and define $H_c^1(S, \Omega_{\mathcal{M}_{g,r}/\mathbf{C}}^{\log})$ to be the hypercohomology of this complex. Thus, we obtain a class

$$\Sigma_{g,r}^c|_S \in H_c^1(S, \Omega_{\mathcal{M}_{g,r}/\mathbf{C}}^{\log})$$

Since ϕ extends to a logarithmic morphism $\phi^{\log} : S^{\log} \rightarrow \overline{\mathcal{M}}_{g,r}^{\log}$, by differentiating, we obtain a morphism of complexes

$$\begin{array}{ccc} \Omega_{\mathcal{M}_{g,r}/\mathbf{C}}^{\log}|_S & \longrightarrow & \mathcal{R}|_B \\ \downarrow d\phi^{\log} & & \downarrow \\ \omega_{S^{\log}/\mathbf{C}} & \xrightarrow{\text{Res}} & \mathcal{O}_B \end{array}$$

In particular, we obtain a natural “trace morphism:”

$$\text{tr} : H_c^1(S, \Omega_{\mathcal{M}_{g,r}/\mathbf{C}}^{\log}) \rightarrow H_c^1(S, \omega_{S^{\log}/\mathbf{C}}) = \mathbf{C}$$

Moreover, since (as we saw in [Mzk1], Chapter I), $\Sigma_{g,r}^c$ is essentially the first Chern class of an ample line bundle, $\text{tr}(\Sigma_{g,r}^c|_S) \in \mathbf{C}$ will simply be the degree of the pull-back of this ample line bundle to S . In particular, as long as ϕ is nonconstant, $\text{tr}(\Sigma_{g,r}^c|_S)$ will be nonzero. Thus, *a fortiori*, $\Sigma_{g,r}^c|_S$ will be nonzero.

Let us summarize what we have done so far. From ϕ , we have defined the cohomology class

$$\Sigma_{g,r}^c|_S \in H_c^1(S, \Omega_{\mathcal{M}_{g,r}/\mathbf{C}}^{\log})$$

simply by using the functoriality of cohomology, i.e., *without using any sort of differentiation*. Moreover, knowing that $\Sigma_{g,r}^c|_S$ is nonzero implies (by Serre duality, taking into account the compact supports) that there exists a nonzero morphism

$$\Omega_{\mathcal{M}_{g,r}/\mathbf{C}}^{\log} \rightarrow \omega_{S^{\log}/\mathbf{C}}$$

Such a nonzero morphism is enough (using stability properties of $\Omega_{\mathcal{M}_{g,r}/\mathbf{C}}^{\log}$ as shown in [EV]) to conclude Vojta conjecture-type inequalities. In fact, in the case, $g = r = 1$, one doesn’t even need to use [EV],

since $\Omega_{\overline{\mathcal{M}}_{g,r}/\mathbb{C}}^{\log}$ will already be a line bundle. Thus, the only time in this discussion that we used differentiation was to *show that the class $\Sigma_{g,r}^c|_S$ is nonzero*.

In some sense, this advance is similar to the introduction of indigenous bundles into the discussion of uniformization theory. That is to say, the K  be uniformization theorem in the complex case in the form of the analytic isomorphism $\tilde{X} \cong H$ (where $\tilde{X} \rightarrow X$ is the universal covering of a compact hyperbolic Riemann surface, and H is the upper half-plane) is an object that cannot even really exist in the p -adic case. However, by thinking of this isomorphism as a special kind of indigenous bundle, one at least obtains an *object* which exists also in the p -adic case, and the “only” problem that remains is to translate the “specialness” of the canonical indigenous bundle into the p -adic context. Similarly, the Kodaira-Spencer morphism, i.e., the derivative of $\phi : S \rightarrow \overline{\mathcal{M}}_{g,r}$ is something which does not even exist in the number field case, whereas various kinds of cohomology groups do exist. Thus, we have gone from working with an object (a derivative) which does not exist in the number field case to an object (a cohomology class) which could conceivably exist also in the number field case, but which has, in the geometric case, a special property which is proven by using techniques (namely, differentiation) proper to the geometric case.

Finally, we remark that it is of the utmost importance in the above discussion to always work with *compactly supported cohomology classes*. Indeed, if one doesn’t do this, then one would have to replace $H^1(S, \omega_{S^{\log}/\mathbb{C}})$ in the above discussion with $H^1(S, \omega_{S^{\log}/\mathbb{C}})$, which is zero if B is nonempty, so the argument given above would completely break down.

 2.3.2. Arithmetic Resolutions of the Schwarz Torsor

Thus, the conclusion of the preceding subsection is that one way to try to produce an analogue of the Kodaira-Spencer morphism in the number field case is to try to construct an *arithmetic* version of the Schwarz torsor (which *a priori*, despite being defined over \mathbb{Z} , is strictly an algebro-geometric object). It is from this point of view that the theory of [Mzk1] and the present book become relevant to the arithmetic Kodaira-Spencer problem. Namely, starting with the *algebraically* nontrivial Schwarz torsor

$$(\overline{\mathcal{S}}_{g,r})_{\mathbb{Z}} \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbb{Z}}$$

we saw in  2 of the Introduction of [Mzk1] that the canonical indigenous bundle defines a canonical *real analytic* section

$$s_H : (\overline{\mathcal{M}}_{g,r})_{\mathbf{C}} \rightarrow (\overline{\mathcal{S}}_{g,r})_{\mathbf{C}}$$

i.e., a *canonical arithmetic trivialization of the Schwarz torsor at the infinite prime*. Moreover, all the other interesting uniformization objects in the complex theory – like, for instance, the Bers embedding – essentially fall naturally out of s_H (see Theorem 2.3 of [Mzk1], Introduction). On the other hand, the classical ordinary theory of [Mzk1] defines a canonical p -adic section

$$(\overline{\mathcal{N}}_{g,r}^{ord})_{\mathbf{Z}_p} \rightarrow (\overline{\mathcal{S}}_{g,r})_{\mathbf{Z}_p}$$

(cf. Theorem 2.8 of [Mzk1], Chapter III), i.e., a *canonical arithmetic trivialization of the Schwarz torsor at the finite prime p* .

At the present time, the author has not obtained any interesting results about the relationship between these various arithmetic trivializations of the Schwarz torsor at different primes. However, it seems not unreasonable to expect that by gluing together these various theories at different primes, one could obtain some sort of global arithmetic Schwarz class over $\mathrm{Spec}(\mathbf{Z}) \cup \infty$ which might be of relevance to solving the arithmetic Kodaira-Spencer problem. We hope to be able to address such issues in greater detail in a future paper.

At any rate, the canonical real analytic section at the infinite prime (along with the Weil-Petersson metric) allow one to think of the Schwarz torsor

$$(\overline{\mathcal{S}}_{g,r})_{\mathbf{Z}} \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbf{Z}}$$

as being equipped with “an integral structure at the infinite prime.” In other words, there is a natural *arithmetic Schwarz torsor* object which exists in the context of Arakelov theory. Thus, one can ask, for instance,

To what extent does this torsor admit trivializations?

For instance, if X is a curve (of type (g, r)) over a number field F (with, say, stable reduction over the ring of integers \mathcal{O}_F of F), then one gets a classifying morphism $\mathrm{Spec}(\mathcal{O}_F) \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbf{Z}}$. Pulling back the Schwarz torsor $(\overline{\mathcal{S}}_{g,r})_{\mathbf{Z}} \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbf{Z}}$ via this classifying morphism thus gives rise to a natural torsor

$$\mathcal{T} \rightarrow \mathrm{Spec}(\mathcal{O}_F)$$

equipped with a canonical trivialization at the infinite primes. Then one can ask,

To what extent do there exist sections of \mathcal{T} over \mathcal{O}_F which are also “integral at infinity,” i.e., whose distance from the canonical trivialization at each infinite prime is ≤ 1 (relative to the Weil-Petersson metric)?

Of course, this question appears to be very difficult in general, but it does at least give a way (cf. the discussion of §2.2.1) of formulating an arithmetic analogue of the notion of “a curve over a function field whose Kodaira-Spencer map vanishes.” Thus, it is likely that a solution to this question would shed light on Vojta’s Conjecture ([Vojta]).

Chapter I: Crys-Stable Bundles

§0. Introduction

The purpose of this Chapter is to define and construct (Theorem 2.7) a fine moduli space for a certain kind of \mathbf{P}^1 -bundle with connection, which we refer to as a *crys-stable bundle*. Such bundles generalize, on the one hand, \mathbf{P}^1 -bundles with connection obtained as projectivizations of rank two stable vector bundles with connection, and, on the other hand, the *indigenous bundles* of [Gunning], [Mzk1]. It turns out that although indigenous bundles are sufficient for the ordinary theory of [Mzk1], for the generalized theory of the present work, one needs to study crys-stable bundles, as well. The fine moduli space of crys-stable bundles forms a crystal valued in the category of smooth algebraic spaces over the moduli stack of stable curves. It is also equipped with a collection of morphisms, which forms a sort of Hodge structure for the crystal. This collection of morphisms defines various substrata of the moduli space (Theorem 3.10). The generic stratum consists of certain stable bundles with connection, while the most special stratum consists of indigenous bundles.

Throughout, we make systematic use of the log structures of [Kato], which enable us to work with a very general sort of “log-curve,” that is, we can handle the case of curves with marked points, as well as singular nodal curves on an equal footing to the smooth case. This leads naturally to the study of degenerations of crys-stable bundles. It turns out that in the study of such degenerations, one inevitably encounters “torally crys-stable bundles” (a slight generalization of crys-stable bundles, in which we allow nonnilpotent monodromy at the marked points). In particular, we will ultimately be interested in the study of degenerations of indigenous bundles, hence of “torally indigenous bundles.” Thus, the last two §’s of this Chapter are devoted to torally indigenous bundles. Although, of course, our ultimate aims are arithmetic, the constructions and definitions of this Chapter are strictly algebraic.

§1. Definition and First Properties

§1.1. Notation Concerning the Underlying Curve

We shall use notation similar to that of Chapter I, §1, of [Mzk1]. Let g and r be nonnegative integers such that $2g - 2 + r \geq 1$. Let $\overline{\mathcal{M}}_{g,r}$ be the *moduli stack of stable curves of genus g , with r marked points*, over \mathbf{Z} , and let $\zeta : \mathcal{C} \rightarrow \overline{\mathcal{M}}_{g,r}$ be the universal curve, with its r marked points $s_1, \dots, s_r : \overline{\mathcal{M}}_{g,r} \rightarrow \mathcal{C}$. Note that $\overline{\mathcal{M}}_{g,r}$ has a natural log structure given by the divisor at infinity. Denote the resulting log stack $\overline{\mathcal{M}}_{g,r}^{\log}$. By taking the divisor which is the union of the s_i and the pull-back of the divisor at infinity of $\overline{\mathcal{M}}_{g,r}$, we get a log structure on \mathcal{C} ; we call the resulting log stack \mathcal{C}^{\log} . The morphism $\zeta : \mathcal{C} \rightarrow \overline{\mathcal{M}}_{g,r}$ extends naturally to a morphism of log stacks $\zeta^{\log} : \mathcal{C}^{\log} \rightarrow \overline{\mathcal{M}}_{g,r}^{\log}$. Now let S be a noetherian scheme. To deal with singular curves, we shall use the theory of log schemes of [Kato]. Thus, we assume that S has a given *fine* ([Kato], §2) log structure, and denote the resulting log scheme by S^{\log} . Let $f^{\log} : X^{\log} \rightarrow S^{\log}$ be a morphism of log schemes.

Definition 1.1. We shall say that $f^{\log} : X^{\log} \rightarrow S^{\log}$ is a *stable log-curve* if there exists a *classifying morphism* $\phi^{\log} : S^{\log} \rightarrow \overline{\mathcal{M}}_{g,r}^{\log}$ such that $X^{\log} \cong S^{\log} \times_{\overline{\mathcal{M}}_{g,r}^{\log}} \mathcal{C}^{\log}$.

We shall denote by $\omega_{X^{\log}/S^{\log}}$ the line bundle of relative logarithmic differentials of $f^{\log} : X^{\log} \rightarrow S^{\log}$, and by $\tau_{X^{\log}/S^{\log}}$ the dual bundle to $\omega_{X^{\log}/S^{\log}}$. We shall denote the marking sections of $f : X \rightarrow S$ by $\sigma_1, \dots, \sigma_r : S \rightarrow X$, and the marking divisor (i.e., the union of the images of the σ_i) by $M_f \subseteq X$.

In what follows, we shall assume that the following assumption is in force:

(*) *The prime 2 is invertible on S .*

§1.2. Definition of a Crys-Stable Bundle

Let $f^{\log} : X^{\log} \rightarrow S^{\log}$ be a stable log-curve. Let $\pi : P \rightarrow X$ be a \mathbf{P}^1 -bundle (in the étale topology on X) with a logarithmic connection ∇_P (relative to f^{\log}).

Let $\tau_{P/X}$ be the relative tangent bundle of P over X . Let us denote by $\text{Ad}(P)$ the vector bundle $\pi_* \tau_{P/X}$ on X . Note that $\text{Ad}(P)$ has a canonical logarithmic connection ∇_{Ad} induced by ∇_P . Moreover,

$\mathrm{Ad}(P)$ also has the structure of a bundle of (simple) *Lie algebras*, whose bracket operation is induced by the bracket operation on sections of $\tau_{P/X}$, regarded as vector fields on P . This bracket action is respected by the connection ∇_{Ad} . Since the Lie algebra structure is semi-simple, it follows that we have a nondegenerate Killing form $\kappa(-, -) : \mathrm{Ad}(P) \otimes_{\mathcal{O}_X} \mathrm{Ad}(P) \rightarrow \mathcal{O}_X$ (given by forming $\frac{1}{4}$ of the trace of the product of the two matrices obtained from the adjoint representation of $\mathrm{Ad}(P)$). The Killing form will also be denoted by $\langle -, - \rangle$. The monodromy operator μ_i of the logarithmic connection ∇_P at σ_i is a section of $\sigma_i^* \mathrm{Ad}(P)$ over S . Thus, $\langle \mu_i, \mu_i \rangle$ forms a section of \mathcal{O}_S over S .

For each $i = 1, \dots, r$, let us fix a *radius* $\rho_i \in \Gamma(S, \mathcal{O}_S)$, which is either 0 or a section of \mathcal{O}_S^\times . Occasionally, it will be convenient to consider the quantities $2\rho_i$, which we shall refer to as *diameters*.

Definition 1.2. We shall say that $(\pi : P \rightarrow X; \nabla_P)$ is a *torally crys-stable bundle* on X^{\log} (of radii $\{\rho_i\}$) if the following conditions are satisfied:

- (1) The monodromy operator $\mu_i \in \Gamma(S, \sigma_i^* \mathrm{Ad}(P))$ of ∇_P at each σ_i satisfies $\langle \mu_i, \mu_i \rangle = 2\rho_i^2$ (for instance, if $\rho_i = 0$, then this condition means that μ_i is square nilpotent).
- (2) The monodromy operator μ_i is nonzero at every point s of S (i.e., the image of μ_i in $(\sigma_i^* \mathrm{Ad}(P)) \otimes_{\mathcal{O}_S} k(s)$ is nonzero).
- (3) Let k be an algebraically closed field; $s : \mathrm{Spec}(k) \rightarrow S$ a morphism; Y an irreducible component of the normalization of $X_s \stackrel{\mathrm{def}}{=} X \times_{S, s} \mathrm{Spec}(k)$; \mathcal{L} a line bundle on Y ; and $\alpha : \mathcal{L} \hookrightarrow \mathrm{Ad}(P)|_Y$ a morphism of vector bundles whose image is stabilized by ∇_{Ad} . Then $\deg(\mathcal{L}) < 0$.
- (4) At every node $\nu \in X$ of the curve $X \rightarrow S$, the monodromy operator $\mu_\nu \in \mathrm{Ad}(P) \otimes_{\mathcal{O}_X} k(\nu)$ is nonzero.

We shall say that $(\pi : P \rightarrow X; \nabla_P)$ is *crys-stable* on X^{\log} if all the radii ρ_i are zero.

Note that it is immediate that the indigenous bundles of [Mzk1], Chapter I, §2, are *crys-stable*. Also, suppose that: $X \rightarrow S$ is a smooth, unmarked curve; \mathcal{E} is a rank two vector bundle with connection on X such that $\mathrm{Ad}(\mathcal{E})$ is a stable vector bundle; and ∇_P is the induced connection on $P \stackrel{\mathrm{def}}{=} \mathbf{P}(\mathcal{E}) \rightarrow X$. Then one checks easily that (P, ∇_P) is *crys-stable*.

Remark. The reason for the name “crys-stable” is that the main conditions on the bundle ((2) to (4)) are “stability conditions” on the

bundle in the following sense: The requirements that the monodromy be nonzero are essentially conditions to the effect that the monodromy operators be “stable as linear operators,” while condition (3) is similar to the usual condition for stability of vector bundles, except that instead of requiring that the degree of *any* subvector bundle satisfy a certain condition, we restrict our attention to subbundles that are stabilized by the connection (hence the “crys-” part of “crys-stable”).

Ultimately, we shall really only be interested in crys-stable bundles (not in the the more general notion of torally crys-stable bundles). The reason, however, that we introduced torally crys-stable bundles here is that in order to prove various properties of crys-stable bundles, it is often useful to apply techniques involving degeneration, and the study of degenerations of crys-stable bundles necessarily involves the introduction of torally crys-stable bundles.

Example 1.3. Note that unlike the case with indigenous bundles, it is not necessarily the case that if $Z^{\log} \rightarrow X^{\log}$ is a log admissible (see, e.g., [Mzk2], §3) morphism of stable log-curves, then the pull-back to Z^{\log} of a crys-stable bundle on X^{\log} remains crys-stable on Z^{\log} . Indeed, one can construct such an example easily as follows: Suppose that $S = \text{Spec}(k)$ (where k is an algebraically closed field), and X is a smooth, unmarked curve of genus g over k . Let $Y \rightarrow X$ be a finite, Galois, étale morphism, with Galois group G , and let $G \rightarrow \text{PGL}_2(k)$ be a representation. This gives an action of G on \mathbf{P}_k^1 . Suppose that G does not fix any unordered pair of (not necessarily distinct) points of \mathbf{P}_k^1 . Let $P \rightarrow X$ be the \mathbf{P}^1 -bundle obtained by taking the quotient of $Y \times_k \mathbf{P}_k^1$ by the diagonal action of G on both factors. The trivial connection on the trivial \mathbf{P}^1 -bundle $Y \times_k \mathbf{P}_k^1 \rightarrow Y$ then induces a connection ∇_P on $P \rightarrow X$, and one sees easily that $(P \rightarrow X; \nabla_P)$ is crys-stable. On the other hand, since the trivial \mathbf{P}^1 -bundle with connection is not crys-stable, it follows that the pull-back of $(P \rightarrow X; \nabla_P)$ to Y is not crys-stable. Ultimately, we shall put other conditions on the crys-stable bundles that we use, so it turns out that this sort of example will not occur in our theory.

§1.3. Isomorphisms

Let us suppose that $(\pi : P \rightarrow X; \nabla_P)$ is a *torally crys-stable bundle* on X^{\log} (of radii $\{\rho_i\}$). Then we have the following result:

Lemma 1.4. *The pair $(\pi : P \rightarrow X; \nabla_P)$ has no nontrivial automorphisms. In particular, we have $(f_{\text{DR}})_* \text{Ad}(P) = 0$, i.e., $(P \rightarrow X; \nabla_P)$ admits no infinitesimal horizontal automorphisms. (See Proposition 1.6 for more on the notation “ $(f_{\text{DR}})_*$.”)*

Proof. The fact that $(f_{\text{DR}})_* \text{Ad}(P) = 0$ follows immediately from assumption (3) of Definition 1.2. Thus, it suffices to prove the result when S is the spectrum of an algebraically closed field k . Let $\alpha : P \rightarrow P$ be an automorphism over X that preserves ∇_P . Note that there exists a finite admissible (as in [Mzk2], §3) covering $Y \rightarrow X$ together with a line bundle \mathcal{L} on $P|_Y \stackrel{\text{def}}{=} P \times_X Y$ such that $\mathcal{L}^{\otimes 2} \cong \tau_{P/X}|_Y$. Since α necessarily preserves $\tau_{P/X}$, it follows that $\alpha^* \mathcal{L} = \mathcal{L} \otimes \pi^* \mathcal{A}$, where \mathcal{A} is a line bundle on Y whose square is trivial. Let $\mathcal{E} = \pi_* \mathcal{L}$. Thus, \mathcal{E} is a rank two line bundle on Y whose projectivization is equal to P , and $\text{Ad}(\mathcal{E})$ (the trace zero endomorphisms of \mathcal{E}) is naturally isomorphic to $\text{Ad}(P)|_Y$. Moreover, the logarithmic connection ∇_P induces a logarithmic connection $\nabla_{\mathcal{E}}$ on \mathcal{E} . Let $\nabla_{\mathcal{A}}$ be the unique connection on \mathcal{A} whose square is the trivial connection on $\mathcal{A}^{\otimes 2} \cong \mathcal{O}_X$.

Now α induces a morphism of vector bundles $\mathcal{E} \rightarrow \mathcal{E} \otimes_{\mathcal{O}_Y} \mathcal{A}$, which is clearly compatible with the connections $\nabla_{\mathcal{E}}$ and $\nabla_{\mathcal{E}} \otimes \nabla_{\mathcal{A}}$. Thus, we obtain a horizontal section $\alpha_{\mathcal{E}} \in \Gamma(Y, \text{End}(\mathcal{E}) \otimes_{\mathcal{O}_Y} \mathcal{A})$. Let $\beta \in \Gamma(Y, \text{Ad}(\mathcal{E}) \otimes_{\mathcal{O}_Y} \mathcal{A})$ be the horizontal section $\alpha_{\mathcal{E}} - \frac{1}{2} \text{tr}(\alpha_{\mathcal{E}})$. Note that β is independent of the choice of \mathcal{L} such that $\mathcal{L}^{\otimes 2} \cong \tau_{P/X}|_Y$. If we regard β as a horizontal morphism of vector bundles $\mathcal{A}^{-1} \rightarrow \text{Ad}(\mathcal{E}) = \text{Ad}(P)|_Y$, then the image of β in $\text{Ad}(P)|_Y$ descends to X . Indeed, this follows from the fact that α descends to X , and β is constructed canonically from α . Thus, if β is nonzero, we obtain a line bundle $\mathcal{B} \subseteq \text{Ad}(P)$ which is stabilized by ∇_{Ad} and of degree ≥ 0 . This contradicts assumption (3) of Definition 1.2, thus completing the proof of the Lemma. \circ

Let $(\rho : Q \rightarrow X; \nabla_Q)$ be another torally crys-stable bundle on X^{\log} (with possibly different radii $\{\rho'_i\}$). Let $\text{Isop}_{P,Q}$ be the functor on S -schemes that assigns to an S -scheme T , the set of isomorphisms $\alpha : Q_T \rightarrow P_T$ of \mathbf{P}^1 -bundles over X_T (where the subscripted “ T ”’s denote pull-back to T from S) that are compatible with the connections ∇_P and ∇_Q .

Proposition 1.5. *The functor $\text{Isop}_{P,Q}$ is representable by a closed subscheme of S .*

Proof. First note that $P \rightarrow X$ can be recovered from the bundle of Lie algebras $\text{Ad}(P)$ by considering the Borel subalgebras of $\text{Ad}(P)$. Thus, instead of considering isomorphisms of torally crys-stable bundles, it suffices to consider isomorphisms of vector bundles $\text{Ad}(P) \rightarrow \text{Ad}(Q)$ that preserve:

- (1) the Lie algebra structures;
- (2) the connections.

It is easy to see that the functor of such isomorphisms is representable by a scheme of finite type over S . This proves that $Isop_{P,Q}$ is representable by such a scheme. By abuse of notation, let us denote this scheme by $Isop_{P,Q}$. Lemma 1.4 implies that $Isop_{P,Q} \rightarrow S$ is a monomorphism (in the category of schemes). Thus, in order to complete the proof of the Proposition, it suffices to prove that $Isop_{P,Q} \rightarrow S$ is proper. For this, we may use the valuative criterion. Thus, we assume that S is the spectrum of a discrete valuation ring A . Let η be the generic point of S . Let us assume that we are given a section of $Isop_{P,Q}$ over η . We may think of this section as a morphism of vector bundles $\beta_\eta : \text{Ad}(P)_\eta \rightarrow \text{Ad}(Q)_\eta$ that respects the various structures listed above.

We claim that β_η is *integral*, i.e., that it extends to a morphism $\beta : \text{Ad}(P) \rightarrow \text{Ad}(Q)$. Indeed, let $x \in A$ be a uniformizer, and let k be the residue field of A . Let us assume that β_η maps $\text{Ad}(P) \subseteq \text{Ad}(P)_\eta$ into $x^{-n} \cdot \text{Ad}(Q)$ (where n is a positive integer), and that this is the minimal n with that property. Let $\beta_k : \text{Ad}(P) \otimes_A k \rightarrow \text{Ad}(Q) \otimes_A k$ be the morphism of vector bundles given by multiplying $\beta_\eta|_{\text{Ad}(P)}$ by x^n and then tensoring over A with k . Since β_η is compatible with the bracket operation, we obtain that $[(x^n \cdot \beta_\eta)(\text{Ad}(P)), (x^n \cdot \beta_\eta)(\text{Ad}(P))] \subseteq x^n \cdot (x^n \cdot \beta_\eta)([\text{Ad}(P), \text{Ad}(P)]) \subseteq x^n \cdot \text{Ad}(Q)$; hence, it follows that β_k maps $\text{Ad}(P)_k \stackrel{\text{def}}{=} \text{Ad}(P) \otimes_A k$ into an abelian subalgebra of $\text{Ad}(Q)_k$. It thus follows that for each connected component Z of the normalization of X_k , there exists a nonzero line bundle $\mathcal{L}_Z \subseteq \text{Ad}(Q)_k|_Z$ on Z which contains the image of $\beta_k|_Z$. Moreover, since β_η is compatible with the connections, it follows that (we may assume that) \mathcal{L}_Z is stabilized by the connection on $\text{Ad}(Q)_k$. If the degree of \mathcal{L} is nonnegative, we thus obtain a contradiction by assumption (3) of Definition 1.2. If the degree of \mathcal{L}_Z is negative, then by considering the dual morphism $\mathcal{L}_Z^\vee \rightarrow \text{Ad}(P)_k|_Z$ (recall that $\text{Ad}(P)$ is self-dual) to $\text{Ad}(P)_k|_Z \rightarrow \mathcal{L}_Z$, we again obtain a contradiction by assumption (3) of Definition 1.2. This completes the proof of the claim that β_η is integral. Thus, β_η extends to a morphism of vector bundles $\beta : \text{Ad}(P) \rightarrow \text{Ad}(Q)$ which is clearly compatible with the connections and Lie algebra structures.

It then follows easily that β is an isomorphism. Indeed, we can apply the argument of the preceding paragraph to β_η^{-1} to conclude that β_η^{-1} maps $\text{Ad}(Q)$ into $\text{Ad}(P)$. This map is clearly inverse to β . This completes the verification of the valuative criterion of properness, and hence the proof of the Proposition. \circ

§1.4. De Rham Cohomology

Let $(\pi : P \rightarrow X; \nabla_P)$ be a *torally crys-stable bundle* on X^{\log} (of radii $\{\rho_i\}$). Since the monodromy at the marked points of radius zero is square nilpotent (and nonzero at every point), over each marked point of radius zero $\sigma_i : S \rightarrow X$, there exists a unique section $q_i : S \rightarrow \sigma_i^* P$

stabilized by the monodromy operator. We shall refer to q_i as the *canonical parabolic structure on the torally crys-stable bundle* $(\pi : P \rightarrow X; \nabla_P)$ at σ_i . On the other hand, at a marked point of nonzero radius, there are precisely two sections $q_i[0], q_i[\infty] : S \rightarrow P$ lying over σ_i that are stabilized by μ_i : one, $q_i[0]$, corresponding to the eigenspace of $\sigma_i^* \text{Ad}(P)$ where μ_i acts with the eigenvalue $2\rho_i$, the other, $q_i[\infty]$, corresponding to the eigenspace where μ_i acts with eigenvalue $-2\rho_i$. We shall refer to the pair $(q_i[0], q_i[\infty])$ as the *canonical toral structure on the torally crys-stable bundle* $(\pi : P \rightarrow X; \nabla_P)$ at σ_i .

Let $\text{Ad}^q(P) \subseteq \text{Ad}(P)$ be the subsheaf of sections s such that $\langle s|_{\sigma_i}, \mu_i \rangle = 0$, for all $i = 1, \dots, r$. If all the radii are zero, i.e., $(P; \nabla_P)$ is *crys-stable* then we define $\text{Ad}^c(P) \subseteq \text{Ad}^q(P)$ to be the subsheaf of sections that vanish to second order (in the relative coordinate for π) at the q_i 's. Next we wish to discuss the de Rham cohomology of the \mathbf{P}^1 -bundle with connection $(P; \nabla_P)$ (taking into account the canonical toral and parabolic structures at the marked points). Note that the exterior differential operator maps $\text{Ad}(P)$ (respectively, $\text{Ad}^c(P)$) into $\text{Ad}^q(P)$ (respectively, $\text{Ad}(P)(-M_f)$). We define the *de Rham cohomology of* $\text{Ad}(P)$ to be the hypercohomology of the complex

$$\text{Ad}(P) \rightarrow \text{Ad}^q(P) \otimes \omega_{X^{\log}/S^{\log}}$$

If all the radii are zero, then we define the *de Rham cohomology with compact supports of* $\text{Ad}(P)$ to be the hypercohomology of the complex

$$\text{Ad}^c(P) \rightarrow \text{Ad}(P) \otimes \omega_{X^{\log}/S^{\log}}(-M_f)$$

Note that there is a natural inclusion from this complex to the complex for the cohomology *without* compact supports. One computes easily that the cohomology of the cokernel of this inclusion is *zero*, i.e., the cohomology *with* compact supports is naturally isomorphic to the cohomology *without* compact supports. Moreover, by Riemann-Roch, Grothendieck-Serre duality, and Lemma 1.4, one obtains the following result:

Proposition 1.6. *Let (P, ∇_P) be a torally crys-stable bundle on an r -pointed stable log-curve $f^{\log} : X^{\log} \rightarrow S^{\log}$ of genus g . Then the de Rham cohomology of $\text{Ad}(P)$ with its natural connection ∇_{Ad} (induced by ∇_P) is as follows:*

- (1) *For cohomology without compact supports, we have $(f_{\text{DR}})_*(\text{Ad}(P)) = \mathbf{R}^2(f_{\text{DR}})_*(\text{Ad}(P)) = 0$; and $\mathbf{R}^1(f_{\text{DR}})_*(\text{Ad}(P))$ is a vector bundle of rank $2(3g - 3 + r)$ on S ;*
- (2) *In the case where all the radii are zero, the cohomology with compact supports is naturally isomorphic to the cohomology without compact supports.*

Moreover, the cup product defines a natural duality between the cohomology without compact supports

$$\mathbf{R}^i(f_{\mathrm{DR}})_*(\mathrm{Ad}(P))$$

and the cohomology with compact supports

$$\mathbf{R}^{2-i}(f_{\mathrm{DR}})_{c,*}(\mathrm{Ad}(P))$$

In particular, (when all the radii are zero) the isomorphism of (2) implies that the cohomology module $\mathbf{R}^1(f_{\mathrm{DR}})_*(\mathrm{Ad}(P))$ is naturally self-dual.

Now let us suppose that we are given a section $h : X \rightarrow P$ of $\pi : P \rightarrow X$ such that the line bundle $h^*\omega_{P/X}$ on X is relatively ample with respect to $X \rightarrow S$, and such that the Kodaira-Spencer morphism $\kappa_h : \tau_{X^{\log}/S^{\log}} \rightarrow h^*\tau_{P/X}$ is an isomorphism at the marked points. Note that condition (3) of Definition 1.2 implies that κ_h is automatically an isomorphism on a dense open set of every fiber of $f : X \rightarrow S$. Let $\mathcal{L} \stackrel{\mathrm{def}}{=} h^*\omega_{P/X}$. Then we obtain a natural inclusion $\mathcal{L} \hookrightarrow \mathrm{Ad}(P)$. Note that the composite of this inclusion with its dual $\mathrm{Ad}(P) \rightarrow \mathcal{L}^\vee$ is zero. Thus, we obtain a decreasing filtration on $\mathrm{Ad}(P)$: $F^{-1}(\mathrm{Ad}(P)) \stackrel{\mathrm{def}}{=} \mathrm{Ad}(P)$; $F^0(\mathrm{Ad}(P)) \stackrel{\mathrm{def}}{=} \mathrm{Ker}(\mathrm{Ad}(P) \rightarrow \mathcal{L}^\vee)$; $F^1(\mathrm{Ad}(P)) \stackrel{\mathrm{def}}{=} \mathrm{Im}(\mathcal{L} \hookrightarrow \mathrm{Ad}(P))$. This filtration induces a natural filtration on the de Rham cohomology of $\mathrm{Ad}(P)$.

Proposition 1.7. *Let \mathcal{E} be the vector bundle $\mathbf{R}^1(f_{\mathrm{DR}})_*(\mathrm{Ad}(P))$ on S . Let $\omega_{X/S}$ be the relative dualizing sheaf of X over S . Then we have the following natural isomorphisms:*

$$(F^{-1}/F^0)(\mathcal{E}) \cong \mathbf{R}^1 f_* \mathcal{L}^{-1}; \quad F^2(\mathcal{E}) \cong f_* \omega_{X/S} \otimes_{\mathcal{O}_X} \mathcal{L}$$

Moreover, the natural isomorphism between cohomology with and cohomology without compact supports induces isomorphisms between the respective F^{-1}/F^0 's and F^2 's.

Proof. The proof is similar to that of Theorem 2.8 of [Mzk1], Chapter I. One simply computes the filtration on the de Rham cohomology by using long exact sequences applied to the filtration on $\mathrm{Ad}(P)$, plus what we know about h : that \mathcal{L} is ample, and that κ_h is generically an isomorphism on every fiber of $X \rightarrow S$, as well as at the marked points.

○

§2. Moduli

§2.1. Boundedness

Let $(\pi : P \rightarrow X; \nabla_P)$ be a torally crys-stable bundle on X^{\log} (of radii $\{\rho_i\}$).

Lemma 2.1. *Let \mathcal{L} be a line bundle on X of the form $\omega_{X^{\log}/S^{\log}} \otimes_{\mathcal{O}_X} \mathcal{M}$, where \mathcal{M} is a line bundle on X which is relatively ample with respect to $f : X \rightarrow S$. Then every morphism of vector bundles $\alpha : \mathcal{L} \rightarrow \mathrm{Ad}(P)$ is zero.*

Proof. Clearly it suffices to prove the Lemma in the case where S is the spectrum of an algebraically closed field k . By replacing X^{\log} by an irreducible component of the normalization of X^{\log} , it follows that we can assume that X is smooth. Thus, we may also assume that the cokernel of α is a vector bundle on X . Consider the morphism $\kappa(\alpha, \alpha) : \mathcal{L}^{\otimes 2} \rightarrow \mathcal{O}_X$. Since \mathcal{L} is ample on X , it follows that $\kappa(\alpha, \alpha) = 0$. Thus, \mathcal{L} arises as the subbundle of $\mathrm{Ad}(P)$ of sections of $\tau_{P/X}$ that vanish to second order at some section $h : X \rightarrow P$ of π . The Kodaira-Spencer morphism at h is a morphism of line bundles $\tau_{X^{\log}/S^{\log}} \rightarrow \mathcal{L}^{-1} = \tau_{X^{\log}/S^{\log}} \otimes \mathcal{M}^{-1}$. Since \mathcal{M} is ample, this Kodaira-Spencer morphism must therefore be zero. But then we obtain a contradiction by assumption (3) of Definition 1.2. \square

Lemma 2.2. *Let $n \geq 4$. Let \mathcal{E} be the vector bundle $\mathrm{Ad}(P) \otimes_{\mathcal{O}_X} (\omega_{X^{\log}/S^{\log}})^{\otimes n}$. Then $\mathbf{R}^1 f_* \mathcal{E} = 0$, and $f^* f_* \mathcal{E} \rightarrow \mathcal{E}$ is surjective.*

Proof. The assertion that $\mathbf{R}^1 f_* \mathcal{E} = 0$ follows from Grothendieck-Serre duality, together with Lemma 2.1. Thus, $f_* \mathcal{E}$ is a vector bundle on S , and commutes with base-change. In particular, it suffices to prove the Lemma over an algebraically closed field k . Then the Lemma follows from considering the long exact cohomology sequence obtained from $0 \rightarrow \mathcal{E} \otimes \mathcal{I} \rightarrow \mathcal{E} \rightarrow \mathcal{E} \otimes (\mathcal{O}_X/\mathcal{I}) \rightarrow 0$ (where $\mathcal{I} \subseteq \mathcal{O}_X$ is an invertible sheaf of ideals such that $\mathcal{I} \otimes_{\mathcal{O}_X} (\omega_{X^{\log}/S^{\log}})^{\otimes 2}$ is ample). \square

§2.2. Definition of Various Functors

As usual, we assume that our set $\rho = \{\rho_i\}$ of radii is fixed throughout. Let $\mathcal{Y}_{X/S}^\rho$ denote the functor on S -schemes defined as follows: For $T \rightarrow S$, we let $\mathcal{Y}_{X/S}^\rho(T)$ be the set of isomorphism classes of torally crys-stable bundles $(P \rightarrow X_T; \nabla_P)$ (of radii ρ) on $X_T^{\log} \stackrel{\mathrm{def}}{=} X^{\log} \times_{S^{\log}} T^{\log}$, where T^{\log} is the log scheme whose underlying scheme is T and whose log structure is obtained by pulling back the log structure on S .

Lemma 2.3. *The functor $\mathcal{Y}_{X/S}^\rho$ is formally smooth over S .*

Proof. Suppose that S is affine, and that $S_0 \subseteq S$ is a closed immersion defined by a square-nilpotent ideal. By Proposition 1.6 (i.e., the fact that the second de Rham cohomology group vanishes), it follows that there are no obstructions to lifting a crys-stable bundle on $X_0^{\log} \stackrel{\text{def}}{=} X^{\log} \times_S S_0$ to a \mathbf{P}^1 -bundle with connection on X^{\log} that satisfies assumption (1) of Definition 1.2. But then assumptions (2), (3) and (4) of Definition 1.2 follow automatically. \circ

Let $\mathcal{R}_{X/S}^\rho$ denote the functor on S -schemes defined as follows: For $T \rightarrow S$, we let $\mathcal{R}_{X/S}^\rho(T)$ be the set of isomorphism classes of torally crys-stable bundles $(P \rightarrow X_T; \nabla_P)$ (of radii ρ) on $X_T^{\log} \stackrel{\text{def}}{=} X^{\log} \times_{S^{\log}} T^{\log}$, together with an isomorphism

$$(f_T)_*(\text{Ad}(P) \otimes_{\mathcal{O}_X} (\omega_{X^{\log}/S^{\log}})^{\otimes 4}) \cong \mathcal{O}_T^N$$

(where $N = 21(g-1) + 12r$), that is, a choice of basis for $(f_T)_*(\text{Ad}(P) \otimes_{\mathcal{O}_X} (\omega_{X^{\log}/S^{\log}})^{\otimes 4})$. Thus, we obtain a natural morphism

$$\mathcal{R}_{X/S}^\rho \rightarrow \mathcal{Y}_{X/S}^\rho$$

of functors. Note that $\mathcal{R}_{X/S}^\rho$ has a natural action of $(\text{GL}_N)_S$ (over $\mathcal{Y}_{X/S}^\rho$) given by acting on the N basis sections of $(f_T)_*(\text{Ad}(P) \otimes_{\mathcal{O}_X} (\omega_{X^{\log}/S^{\log}})^{\otimes 4})$. Moreover, this natural action defines a natural isomorphism

$$(\text{GL}_N)_S \times_S \mathcal{R}_{X/S}^\rho \cong \mathcal{R}_{X/S}^\rho \times_{\mathcal{Y}_{X/S}^\rho} \mathcal{R}_{X/S}^\rho$$

functors over S . Finally, let us observe that this natural action defines a morphism

$$(\text{GL}_N)_S \times_S \mathcal{R}_{X/S}^\rho \rightarrow \mathcal{R}_{X/S}^\rho \times_S \mathcal{R}_{X/S}^\rho$$

(given on T -valued points by $(g, \rho) \mapsto (g \cdot \rho, \rho)$) which is a monomorphism (by Lemma 1.4) and satisfies the valuative criterion of properness (by Proposition 1.5).

Lemma 2.4. *Suppose that $\mathcal{R}_{X/S}^\rho$ is representable by an S -scheme of finite type. Then $\mathcal{Y}_{X/S}^\rho$ is representable by a smooth separated algebraic space of finite type over S whose relative dimension over S is $2(3g-3+r)$.*

Proof. Let us assume that $\mathcal{R}_{X/S}^\rho$ is representable by an S -scheme of finite type. Then the preceding observation implies that the morphism

$$(\mathrm{GL}_N)_S \times_S \mathcal{R}_{X/S}^\rho \rightarrow \mathcal{R}_{X/S}^\rho \times_S \mathcal{R}_{X/S}^\rho$$

is a closed immersion. It is then standard that (by forming “étale slices” of $\mathcal{R}_{X/S}^\rho \rightarrow \mathcal{Y}_{X/S}^\rho$) there exists a separated algebraic space representing $\mathcal{Y}_{X/S}^\rho$ (which, by abuse of notation, we also denote by $\mathcal{Y}_{X/S}^\rho$) such that $\mathcal{R}_{X/S}^\rho \rightarrow \mathcal{Y}_{X/S}^\rho$ is a torsor over $(\mathrm{GL}_N)_S$. The assertion on the relative dimension over S follows from Proposition 1.5 (the dimension of the first de Rham cohomology group). \circ

§2.3. Representability

The purpose of this subsection is to prove that $\mathcal{R}_{X/S}^\rho$ is representable by an S -scheme of finite type. Let $G \rightarrow S$ be the Grassmannian of quotients of \mathcal{O}_S^N (where $N = 21(g-1) + 12r$) of rank three. Let $H \rightarrow S$ be the Hilbert scheme of S -morphisms $X \rightarrow G$ such that the pull-back to X of the tautological quotient vector bundle of rank three has relative degree (with respect to $f : X \rightarrow S$) equal to $12(2g-2+r)$. Thus, $H \rightarrow S$ is an S -scheme of finite type. Let $\phi : X_H \stackrel{\mathrm{def}}{=} X \times_S H \rightarrow G_H \stackrel{\mathrm{def}}{=} G \times_S H$ be the tautological morphism over H . Let \mathcal{F} be the pull-back to X_H via ϕ of the tautological quotient vector bundle of rank three on G_H . Let $\mathcal{A} = \mathcal{F} \otimes_{\mathcal{O}_X} (\omega_{X^{\log}/S^{\log}})^{\otimes -4}$. Let $H_1 \rightarrow H$ denote the H -scheme of surjective vector bundle morphisms $\mathcal{A} \otimes_{\mathcal{O}_{X_H}} \mathcal{A} \rightarrow \mathcal{A}$ that define a bracket operation for a Lie algebra structure on \mathcal{A} . Let $Z \rightarrow H_1$ denote the H_1 -scheme parametrizing logarithmic connections $\nabla_{\mathcal{A}}$ on \mathcal{A} (which are most easily parametrized when considered as isomorphisms of the two pull-backs of \mathcal{A} to the first infinitesimal neighborhood of the diagonal in $X^{\log} \times_{S^{\log}} X^{\log}$) that are compatible with the bracket operation $\mathcal{A} \otimes \mathcal{A} \rightarrow \mathcal{A}$, have nonzero monodromy at the nodes (condition (4) of Definition 1.2), and whose monodromy operators at the marked points satisfy conditions (1) and (2) of Definition 1.3.

Let us summarize what we have done so far. We have an S -scheme $Z \rightarrow S$, together with a rank three vector bundle \mathcal{A} on X_Z and a logarithmic connection $\nabla_{\mathcal{A}}$ (with respect to $f_Z^{\log} : X_Z^{\log} \rightarrow S_Z^{\log}$) on \mathcal{A} whose monodromy operators at the marked points and nodes satisfy conditions (1), (2) and (4) of Definition 1.3. Moreover, we have a bracket operation $\mathcal{A} \otimes \mathcal{A} \rightarrow \mathcal{A}$ which gives \mathcal{A} the structure of a bundle of semi-simple (since the bracket operation is surjective) Lie algebras. Thus, we have a nondegenerate Killing form $\kappa(-, -) : \mathcal{A} \otimes_{\mathcal{O}_X} \mathcal{A} \rightarrow \mathcal{O}_X$ (given locally as $\frac{1}{4}$ of the trace of the product of the two matrices obtained from the adjoint representation of $\mathrm{Ad}(P)$). In particular, \mathcal{A} is self-dual. By considering line bundle quotients $\lambda : \mathcal{A} \rightarrow \mathcal{L}$ whose

dual $\lambda^\vee : \mathcal{L}^\vee \rightarrow \mathcal{A}^\vee \cong \mathcal{A}$ satisfies $\kappa(\lambda^\vee, \lambda^\vee) = 0$, we thus obtain (in a natural fashion) a \mathbf{P}^1 -bundle $\pi : P \rightarrow X_Z$ equipped with a logarithmic connection ∇_P . Moreover, we have a natural horizontal isomorphism of $\mathrm{Ad}(P)$ with \mathcal{A} (as follows from the semi-simplicity of \mathcal{A}). Now we claim that:

(*)The set U of points $z \in Z$ at which $(\mathrm{Ad}(P), \nabla_{\mathrm{Ad}})$ satisfies condition (3) of Definition 1.2 is open in Z .

Suppose we know that the claim is true. Then it follows immediately that U represents the functor $\mathcal{R}_{X/S}^\rho$ defined in the preceding subsection.

Let us prove the above claim (*). Let $F = Z - U$. Then we have the following:

Lemma 2.5. *The set F is closed under specialization.*

Proof. This follows immediately by applying the valuative criterion: Thus, we assume that $Z = S$ is the spectrum of a discrete valuation ring A , with some log structure. Also, we have a log-curve $X^{\mathrm{log}} \rightarrow S^{\mathrm{log}}$, and a bundle with connection (P, ∇_P) which is crys-stable over s , the special point of S . It is not difficult to reduce (by replacing X_η by one of its irreducible components) to the case when X_η is smooth. Thus, we assume (for convenience) that X_η is smooth, and that the log structure on S is defined by the special point. Suppose that we have a horizontal section $\psi_\eta : X_\eta \rightarrow \mathbf{P}(\mathrm{Ad}(P))$ whose canonical height (i.e., the degree of the pull-back to X_η of the natural $\mathcal{O}(1)$ on $\mathbf{P}(\mathrm{Ad}(P))$) is ≤ 0 . Then there exists a normal, S -flat Y , together with a morphism $Y \rightarrow X$ which is an isomorphism over η such that ψ_η extends to a (horizontal) section $\psi_Y : Y \rightarrow \mathbf{P}(\mathrm{Ad}(P))$ over Y . Note that the canonical height of $(\psi_Y)_s$ is equal to that of $(\psi_Y)_\eta = \psi_\eta$ (since Y is S -flat), hence ≤ 0 . Let $Z \subseteq Y_s$ be an *exceptional* irreducible component, i.e., one that maps to a point of X_s . Then, since $\mathbf{P}(\mathrm{Ad}(P))$ restricts to a trivial projective bundle on Z , $\psi_Z \stackrel{\mathrm{def}}{=} (\psi_Y)|_Z$ has height ≥ 0 . Since, however, the total height of $(\psi_Y)_s$ is ≤ 0 , it follows that there exists a nonexceptional component Z of Y_s such that ψ_Z has height ≤ 0 . Thus, $\psi_Z : Z \rightarrow \mathbf{P}(\mathrm{Ad}(P))$ is a horizontal section of height ≤ 0 , which contradicts the assumption that (P, ∇_P) is crys-stable on X_s . This completes the proof. \circ

Thus, in order to prove that U is open, it suffices to prove the following:

Lemma 2.6. *The set U is constructible.*

Proof. For the purposes of proving this Lemma, we may assume that $S = Z$. By Lemma 2.5 (and a standard criterion for constructibility - see [Mats]), it suffices to prove the following result:

(†) Assume that S is an integral scheme, and that the generic point η of S belongs to U . Then U contains a nonempty open subset of S .

In the course of proving (†), we are always free to replace S by an open subset of S . Also, by substituting for X_η a connected component of the normalization of X_η (and then replacing S by an open subset of S), it is easy to see that we may assume that $X \rightarrow S$ is *smooth*. Let $Q \rightarrow S$ be the scheme of finite type parametrizing pairs consisting of a line bundle \mathcal{L} on X of relative degree zero over S , together with a nonzero morphism $\alpha : \mathcal{L} \rightarrow \text{Ad}(P)$. Let $W \subseteq Q$ denote the closed subscheme where the image of α is stabilized by ∇_{Ad} . Then note that the subset $F \subseteq S$ is precisely the image of the morphism $W \rightarrow S$. Since this morphism is of finite type, it is well-known ([Mats]) that its image is constructible. Thus, if $\eta \in U = S - F$, it follows that there exists an open neighborhood of η in U . \circ

This completes the proof of the claim (*). It thus, follows that U represents the functor $\mathcal{R}_{X/S}^\rho$ defined in the preceding subsection. Thus, we see (by Lemma 2.4) that we have proven the following result, which is the main result of this Chapter:

Theorem 2.7. *The functor $\mathcal{Y}_{X/S}^\rho$ of torally crys-stable bundles on X^{\log} of radii ρ is representable by a smooth, separated algebraic space of finite type over S whose relative dimension over S is $2(3g - 3 + r)$. Moreover, the formation of $\mathcal{Y}_{X/S}^\rho$ commutes with arbitrary base-change $T^{\log} \rightarrow S^{\log}$.*

Finally, let us consider the universal case. Let A be a $\mathbf{Z}[\frac{1}{2}]$ -algebra. Let us assume that our radii ρ are elements of A . Let $(\overline{\mathcal{M}}_{g,r}^{\log})_A$ denote the log stack of r -pointed stable curves of genus g over A , and let $\zeta_A^{\log} : \mathcal{C}_A^{\log} \rightarrow (\overline{\mathcal{M}}_{g,r}^{\log})_A$ denote the universal log-curve. Then by taking $f^{\log} = \zeta_A^{\log}$; $S^{\log} = (\overline{\mathcal{M}}_{g,r}^{\log})_A$; $X^{\log} = \mathcal{C}_A^{\log}$; we obtain (from Theorem 2.7) a relative smooth, quasi-compact algebraic space

$$(\overline{\mathcal{Y}}_{g,r}^\rho)_A \rightarrow (\overline{\mathcal{M}}_{g,r}^{\log})_A$$

of relative dimension $2(3g - 3 + r)$ over $(\overline{\mathcal{M}}_{g,r}^{\log})_A$. Let $(\mathcal{Y}_{g,r}^\rho)_A \subseteq (\overline{\mathcal{Y}}_{g,r}^\rho)_A$ be the open substack where the curve $X \rightarrow S$ is *smooth*. We equip $(\overline{\mathcal{Y}}_{g,r}^\rho)_A$ with the log structure obtained by pulling back the log structure of $(\overline{\mathcal{M}}_{g,r}^{\log})_A$, and call the resulting log stack $(\overline{\mathcal{Y}}_{g,r}^\rho)_A^{\log}$. Thus, we have a morphism of log stacks

$$(\overline{\mathcal{Y}}_{g,r}^\rho)_A^{\log} \rightarrow (\overline{\mathcal{M}}_{g,r}^{\log})_A$$

§2.4. Radimmersions

Before proceeding, we take some time out here to develop a little basic algebraic geometry which will of use in what follows. The purpose of this algebraic geometry is to give a criterion for a morphism of schemes (or algebraic spaces) to be an immersion that depends only on the functors that the schemes represent. In this subsection, let S be a noetherian scheme. Let X and Y be S -schemes of finite type. Then recall that a morphism $f : X \rightarrow Y$ is said to be *radicial* if every geometric fiber $X_{\bar{y}} \rightarrow \bar{y}$ (where $\bar{y} \rightarrow Y$ is a point valued in an algebraically closed field) has at most one (geometric) point. A morphism $f : X \rightarrow Y$ is an *immersion* if it factors as $X \hookrightarrow U \hookrightarrow Y$, where $X \hookrightarrow U$ is a closed immersion, and $U \hookrightarrow Y$ is an open immersion. Now we make the following

Definition 2.8. A morphism $f : X \rightarrow Y$ is said to be a *radimmersion* if it factors as $X \rightarrow U \hookrightarrow Y$, where $X \rightarrow U$ is a finite radicial morphism, and $U \hookrightarrow Y$ is an open immersion.

Let $f : X \rightarrow Y$ be a morphism. Recall that a *trait* T is, by definition, the spectrum of a discrete valuation ring. If T is a trait, we will denote its generic point by η and its closed point by t . Now consider the following *valuative criterion* concerning f :

(*) Suppose that we are given a trait T , and a pair of commutative diagrams:

$$\begin{array}{ccc} \eta & \longrightarrow & X \\ \downarrow & & \downarrow \\ T & \longrightarrow & Y \end{array} \qquad \begin{array}{ccc} t & \longrightarrow & X \\ \downarrow & & \downarrow \\ T & \longrightarrow & Y \end{array}$$

where the bottom morphisms in the two diagrams are the same; the horizontal morphisms on the left in both diagrams are the natural ones; and the horizontal morphisms on the right in both diagrams are equal to f . Then there exists a *unique* morphism $T \rightarrow X$ that (when inserted in the above two diagrams as a diagonal arrow originating in the lower left corner and terminating at the upper right corner) makes the above diagrams commute.

Now we have the following results:

Lemma 2.9. Suppose that f satisfies (*). Then f is radicial.

Proof. Clearly, $(*)$ is preserved under base change. Thus, we reduce immediately to the case where $Y = \operatorname{Spec}(k)$, and k is an algebraically closed field. Suppose that X has two distinct geometric points $\xi, \zeta \in X(k)$. Then take $T = \operatorname{Spec}(k[[x]])$ (where x is an indeterminate), and let $T \rightarrow Y$ be the morphism arising from the k -algebra structure of $k[[x]]$. Let $\eta \rightarrow X$ be the morphism given by composing $\eta \rightarrow T \rightarrow Y = \operatorname{Spec}(k)$ with $\xi : \operatorname{Spec}(k) \rightarrow X$. Let $t = \operatorname{Spec}(k) \rightarrow X$ be ζ . Then applying $(*)$ clearly results in a contradiction. \circ

Lemma 2.10. *Suppose that f satisfies $(*)$. Then f is separated.*

Proof. Indeed, suppose that there exist two distinct $\alpha, \beta : T \rightarrow X$ that make the first diagram in $(*)$ commute. Then observe that Lemma 2.9 implies that $\alpha|_t = \beta|_t$. Thus, the fact that $\alpha = \beta$ follows from the uniqueness assertion inherent in $(*)$. \circ

Lemma 2.11. *Suppose that f is a radimmersion. Then f satisfies $(*)$.*

Proof. Write $f : X \rightarrow U \hookrightarrow Y$, where $X \rightarrow U$ is finite radicial, and $U \subseteq Y$ is an open subscheme. Suppose that we are in the situation of $(*)$. Then the morphism $T \rightarrow Y$ maps t into U ; since U is open, it thus follows that $T \rightarrow Y$ factors through U . Thus, one can replace Y by U , and assume that f is finite radicial. Then f is proper, so it follows that there exists a unique morphism $\alpha : T \rightarrow X$ that makes the first diagram commute. Moreover, the second diagram also commutes since $X \rightarrow Y$ (being radicial) is injective on points valued in a field. This completes the proof. \circ

Theorem 2.12. *A morphism $f : X \rightarrow Y$ is a radimmersion if and only if f satisfies the valuative criterion $(*)$.*

Proof. It suffices to assume that f satisfies $(*)$ and prove that f is a radimmersion. One knows already that f is radicial. Let us consider first the case where $Y = \operatorname{Spec}(R)$; R is a henselian local ring; and there exists a point $\xi \in X$ that is mapped by f to the closed point of Y . Then since f is quasi-finite, it follows from Zariski's main theorem that X is a disjoint union of open subsets U_i (for $i = 1, \dots, r$) of schemes $\operatorname{Spec}(C_i)$, where C_i is a finite local R -algebra. Without loss of generality, we may assume that $\xi \in U_1$. Thus, $U_1 = \operatorname{Spec}(C_1) \subseteq X$, so the R -algebra C_1 is radicial. I claim that $r = 1$. Indeed, suppose that $u \in U_i$, where $i > 1$. Let $z = f(u) \in Y$. Let $Z \subseteq Y$ be the integral closed subscheme whose generic point is z . Write $Z = \operatorname{Spec}(B)$. Thus, $B \subseteq k(z) \subseteq k(u)$. Now by basic commutative algebra, there exists a discrete valuation ring A whose quotient field K contains $k(u)$ and such that $B \subseteq A$ is a local morphism. By applying $(*)$ to the obvious morphisms, it thus

follows that $T \stackrel{\text{def}}{=} \text{Spec}(A) \rightarrow \text{Spec}(B) = Z \subseteq Y$ factors through X . Denote the resulting morphism by $\alpha : T \rightarrow X$. Since $\alpha(\eta) \in U_i$, it follows that $\alpha(t) \in U_i$. But $f(\alpha(t))$ is the closed point of y , so, since f is radical, $\alpha(t) = \xi \in U_1$, which is absurd. This completes the proof of the claim that $r = 1$. Thus, in summary, we see that in the case under consideration, $f : X = U_1 = \text{Spec}(C_1) \rightarrow Y$ is finite radical.

Now let us return to the case of arbitrary $f : X \rightarrow Y$. For each $x \in X$, let $y = f(x)$, and let Y_y^h be the henselization of Y at y . Let $X_y^h = X \times_Y Y_y^h$. By what we did in the preceding paragraph, we know that $X_y^h \rightarrow Y_y^h$ is finite radical. Let Y_y be the Zariski localization of Y at y ; let $X_y = X \times_Y Y_y$. Then by descent, it follows that $X_y \rightarrow Y_y$ is finite radical. By “spreading out,” it thus follows that there exists an open neighborhood U_x of y in Y such that over U_x , f is finite radical. Let $U \subseteq Y$ be the union of the U_x over all $x \in X$. Then it follows immediately that over U , f is finite radical. On the other hand, by definition, $f : X \rightarrow Y$ factors through U . This completes the proof. \circ

Corollary 2.13. *Suppose that $f : X \rightarrow Y$ is a monomorphism that satisfies $(*)$. Then $f : X \rightarrow Y$ is an immersion.*

Proof. This follows from Theorem 2.12 and the fact that a finite monomorphism is (by Nakayama’s Lemma) a closed immersion. \circ

§3. Further Structure

§3.1. Crystal in Algebraic Spaces

Let $f^{\log} : X^{\log} \rightarrow S^{\log}$ be a stable log-curve. Let $\text{Crys}(S^{\log})$ be the étale, nilpotent crystalline site of the log scheme S^{\log} : that is, objects consists of data $(U \rightarrow S; U^{\log} \hookrightarrow T^{\log}; \gamma)$, where $U \rightarrow S$ is an étale morphism; $U^{\log} \hookrightarrow T^{\log}$ is an *exact closed immersion of log schemes* (as in [Kato], §3.1); and γ is a *PD-nilpotent divided power structure* on the ideal defining the closed immersion $U \hookrightarrow T$. See [Kato], §5, for more details.

Proposition 3.1. *The algebraic space $\mathcal{Y}_{X/S}^p \rightarrow S^{\log}$ arises as the value on the object $(S = S; S^{\log} \hookrightarrow S^{\log})$ of $\text{Crys}(S^{\log})$ of a “natural” crystal on $\text{Crys}(S^{\log})$ valued in the category of separated algebraic spaces. Here, by “natural” we mean that the formation of this crystal is functorial with respect to base-change $T^{\log} \rightarrow S^{\log}$. In particular, the relative algebraic space $\mathcal{Y}_{X/S}^p \rightarrow S$ (regarded as a fiber bundle over S whose fibers are algebraic spaces) admits a natural logarithmic (“Gauss-Manin”) connection on S^{\log} .*

Proof. Let $S^{\log} \hookrightarrow T^{\log}$ be a PD -nilpotent PD -thickening. Suppose that we are given two stable log-curves $Y^{\log} \rightarrow T^{\log}$ and $Z^{\log} \rightarrow T^{\log}$ whose restrictions to S^{\log} coincide with $X^{\log} \rightarrow S^{\log}$. Then it suffices to construct a natural bijection between torally crys-stable bundles on Y^{\log} (of radii ρ) and torally crys-stable bundles on Z^{\log} (of radii ρ) with the property that torally crys-stable bundles on Y^{\log} and Z^{\log} (of radii ρ) that correspond under this bijection induce the same bundle when restricted to X^{\log} . We do this as follows: Let $(P \rightarrow Y, \nabla_P)$ be a torally crys-stable bundle on Y^{\log} of radii ρ . Then $(P \rightarrow Y, \nabla_P)$ defines a crystal on the relative crystalline site $\text{Crys}(Y^{\log}/T^{\log})$ (of PD -nilpotent PD -thickenings) valued in the category of \mathbf{P}^1 -bundles. Since crystals on $\text{Crys}(Y^{\log}/T^{\log})$ (respectively, $\text{Crys}(Z^{\log}/T^{\log})$) and $\text{Crys}(X^{\log}/T^{\log})$ are equivalent, we thus see that we obtain a crystal $(Q \rightarrow Z, \nabla_Q)$ on $\text{Crys}(Z^{\log}/T^{\log})$. The natural bijective correspondence is then given by $(P \rightarrow Y, \nabla_P) \mapsto (Q \rightarrow Z, \nabla_Q)$. \circ

§3.2. Hodge Morphisms

Let $f^{\log} : X^{\log} \rightarrow S^{\log}$ be a stable log-curve (of genus g , with r marked points). Let $\mathcal{Y}_{X/S}^{\rho} \rightarrow S$ be the algebraic space of torally crys-stable bundles of radii ρ , as given by Theorem 2.7. Let

$$\chi \stackrel{\text{def}}{=} \frac{1}{2}(2g - 2 + r)$$

Let $l \in \frac{1}{2}\mathbf{Z}$ be a positive half-integer. Then we make the following:

Definition 3.2. We define $\mathcal{Y}_{X/S}^{\rho;l} \rightarrow S$ to be the functor on S -schemes that assigns to $T \rightarrow S$ the set of isomorphism classes of pairs consisting of: (I.) a torally crys-stable bundle $(\pi : P \rightarrow X_T, \nabla_P)$ on $X_T^{\log} \rightarrow T^{\log}$ of radii ρ ; (II.) a section $h : X_T \rightarrow P$ of $\pi : P \rightarrow X_T$ such that the following properties are satisfied:

- (1) the Kodaira-Spencer morphism $\kappa_h : \tau_{X^{\log}/S^{\log}}|_T \rightarrow h^*(\tau_{P/X_T})$ of h is an isomorphism at the nodes and marked points of $X_T \rightarrow T$;
- (2) the line bundle $h^*\omega_{P/X_T}$ on X_T is relatively ample for $X_T \rightarrow T$, and has relative degree $\deg_{X_T/T}(h^*\omega_{P/X_T})$ equal to $2l$.

In general, we shall refer to the number $\frac{1}{2} \deg_{X_T/T}(h^*\tau_{P/X_T})$ as the *canonical height of the section h* . Moreover, we define $\mathcal{Y}_{X/S}^{\rho;0} \subseteq \mathcal{Y}_{X/S}^{\rho}$ to be the subfunctor whose T -valued points are the torally crys-stable bundles $(\pi : P \rightarrow X_T, \nabla_P)$ on X_T^{\log} such that for any geometric point $t : \text{Spec}(k) \rightarrow T$ (where k is an algebraically closed field), every section $h : Y \rightarrow P|_Y$ over an connected component Y of the

normalization of X_s has nonnegative canonical height. We shall call an S -valued point of $\mathcal{Y}_{X/S}^{\rho;l}$ a *torally crys-stable bundle on X^{\log} of level l and radii ρ* , and we shall refer to the section $h : X \rightarrow P$ as the *Hodge section of the torally crys-stable bundle in question*. (Note that the use of the definite article “the” (preceding “Hodge section”) is justified by Lemma 3.5 below.)

Note that the Kodaira-Spencer morphism $\kappa_h : \tau_{X^{\log}/S^{\log}}|_T \rightarrow h^* \tau_{P/X_T}$ is nonzero on an open dense subset of every fiber of $X_T \rightarrow T$. Indeed, this follows immediately from condition (3) of Definition 1.2 and condition (2) of Definition 3.2. Note also that we have a natural morphism $\mathcal{Y}_{X/S}^{\rho;l} \rightarrow \mathcal{Y}_{X/S}^{\rho}$ given by forgetting the section $h : X_T \rightarrow P$.

Proposition 3.3. *This morphism $\mathcal{Y}_{X/S}^{\rho;l} \rightarrow \mathcal{Y}_{X/S}^{\rho}$ is representable by a relative (over $\mathcal{Y}_{X/S}^{\rho}$) scheme of finite type, for all $l \geq 0$. Moreover, when $l = 0$, this morphism is an open immersion.*

Proof. For $l > 0$, the Proposition follows immediately from the theory of Hilbert schemes. When $l = 0$, one checks easily that the set of points of $\mathcal{Y}_{X/S}^{\rho}$ in the image $\mathcal{Y}_{X/S}^0$ is constructible (cf. the proof of Lemma 2.6) and closed under generization (cf. the proof of Lemma 2.5). \circ

As usual, by abuse of notation, we denote this relative scheme by $\mathcal{Y}_{X/S}^{\rho;l} \rightarrow \mathcal{Y}_{X/S}^{\rho}$. Thus, $\mathcal{Y}_{X/S}^{\rho;l}$ is an algebraic space over S .

Lemma 3.4. *When $l = \chi$, $\mathcal{Y}_{X/S}^{\rho;l} \rightarrow S$ is precisely the moduli space of isomorphism classes of torally indigenous bundles of radii ρ (see §4 for more details). When $l > \chi$, $\mathcal{Y}_{X/S}^{\rho;l}$ is the empty scheme.*

Proof. For simplicity, instead of working with arbitrary T -valued points, we assume $T = S$. Suppose that we have a section $h : P \rightarrow X$ of canonical height $\leq -\chi$. Then since κ_h is generically nonzero on every fiber of $X \rightarrow S$, it follows immediately from degree considerations that it must be an isomorphism. Thus, the canonical height of h is precisely $-\chi$. \circ

Let Π be the disjoint union of the $\mathcal{Y}_{X/S}^{\rho;j}$, as j ranges over all non-negative half-integers. Thus, Π is an algebraic space of finite type over S , equipped with a natural morphism $\Pi \rightarrow \mathcal{Y}_{X/S}^{\rho}$.

Lemma 3.5. *The morphism $\Pi \rightarrow \mathcal{Y}_{X/S}^{\rho}$ is a monomorphism.*

Proof. Suppose that (P, ∇_P) is a torally crys-stable bundle (of radii ρ) on X^{\log} and $h_1, h_2 : X \rightarrow P$ are two sections of canonical heights $-l_1$ and $-l_2$, respectively. (Here we assume $l_1, l_2 > 0$.) Then h_1 (respectively, h_2) corresponds to a line bundle $\mathcal{L}_1 \subseteq \text{Ad}(P)$ (respectively, $\mathcal{L}_2 \subseteq \text{Ad}(P)$) of degree $2l_1$ (respectively, $2l_2$), where, as usual, for $n = 1, 2$, $\mathcal{L}_n \subseteq \text{Ad}(P) = \pi_* \tau_{P/X}$ is the set of sections of $\tau_{P/X}$ that vanish to second order at h_n . By degree considerations, the composite $\mathcal{L}_1 \rightarrow \mathcal{L}_2^\vee$ of the inclusion $\mathcal{L}_1 \subseteq \text{Ad}(P)$ with the dual morphism $\text{Ad}(P) \rightarrow \mathcal{L}_2^\vee$ is identically zero. Thus, $\text{Ad}(P) \rightarrow \mathcal{L}_2^\vee$ factors through $\text{Ad}(P)/\mathcal{L}_1$. On the other hand, there is a natural injection $\mathcal{O}_X \hookrightarrow \text{Ad}(P)/\mathcal{L}_1$ (whose cokernel is \mathcal{L}_1^\vee). Again, by degree considerations the composite $\mathcal{O}_X \hookrightarrow \text{Ad}(P)/\mathcal{L}_1 \rightarrow \mathcal{L}_2^\vee$ is zero. Now, putting this all together, if one thinks of the evaluation morphism $\pi^* \text{Ad}(P) \rightarrow \tau_{P/X}$ as defining an embedding (over X) of P into $\mathbf{P}(\text{Ad}(P))$, we see that we have proven that $\text{Im}(h_2) = \text{Im}(h_1)$, i.e., $h_1 = h_2$, as desired. \circ

Now suppose that we are given an S -valued point of $\mathcal{Y}_{X/S}^{\rho;l}$, i.e., a torally crys-stable bundle $(\pi : P \rightarrow X; \nabla_P)$ of radii ρ on X^{\log} , together with a section $h : X \rightarrow P$ satisfying the properties listed in Definition 3.2. Let $D_\kappa \subseteq X$ be the zero locus of the Kodaira-Spencer morphism κ_h of h .

Definition 3.6. We shall refer to the divisor $D_\kappa \subseteq X$ as the *Kodaira-Spencer locus* of (P, ∇_P) .

Definition 3.7. Let $l > 0$. Let $\overline{\mathcal{D}}_{g,r}^l \rightarrow \overline{\mathcal{M}}_{g,r}$ denote the stack of pairs consisting of an r -pointed stable curve $f : X \rightarrow S$ of genus g , together with an effective relative divisor $\Delta_\kappa \subseteq X$ (over S) such that:

- (1) Δ_κ avoids the nodes and marked points of X ;
- (2) the line bundle $\omega_{X^{\log}/S^{\log}}(-\Delta_\kappa)$ on X is relatively ample over S of relative degree $2l$.

We shall call divisors Δ_κ satisfying the above two conditions *l-balanced*.

It is easy to see that $\overline{\mathcal{D}}_{g,r}^l \rightarrow \overline{\mathcal{M}}_{g,r}$ is relatively representable by a smooth quasi-projective relative scheme of relative dimension $2\chi - 2l$. Note that the Kodaira-Spencer divisor $D_\kappa \subseteq X$ defines an S -valued point of $\mathcal{D}_{X/S}^l \stackrel{\text{def}}{=} \overline{\mathcal{D}}_{g,r}^l \times_{\overline{\mathcal{M}}_{g,r}} S$ (where the morphism $S \rightarrow \overline{\mathcal{M}}_{g,r}$ appearing in the fiber product is the classifying morphism for the log-curve $X^{\log} \rightarrow S^{\log}$). Indeed, this follows immediately from Definition 3.2. Thus, we obtain a natural morphism

$$\Delta_{X/S}^l : \mathcal{Y}_{X/S}^{\rho;l} \rightarrow \mathcal{D}_{X/S}^l$$

of algebraic spaces over S . In some sense, this morphism may be regarded as an auxiliary portion of the “Hodge structure” on the crystal defined by the space of torally crys-stable bundles. Indeed, one has the following:

Lemma 3.8. *The morphism $\Delta_{X/S}^l$ is smooth of relative dimension $3g - 3 + r$. Moreover, at an S -valued point of $\mathcal{Y}_{X/S}^{\rho;l}$ corresponding to a pair $(P \rightarrow X, \nabla_P)$ on X^{\log} , the surjection on relative tangent bundles over S induced by $\Delta_{X/S}^l$ at this point may be identified with*

$$F^0(\mathbf{R}^1(f_{\mathrm{DR}*}(\mathrm{Ad}(P))) \rightarrow (F^0/F^1)(\mathbf{R}^1(f_{\mathrm{DR}*}(\mathrm{Ad}(P))))$$

(where “ F^0 ” and “ F^1 ” arise from the “Hodge filtration” as discussed in §1.4). In particular, the morphism $\mathcal{Y}_{X/S}^{\rho;l} \rightarrow S$ is smooth of relative dimension $2(3g - 3 + r) - 2l - (g - 1)$ (if $0 < l \leq \chi$).

Proof. We consider obstructions to infinitesimal deformations (as in Lemma 2.3). First of all, the infinitesimal deformations of the torally crys-stable bundle (P, ∇_P) are given by the bundle

$$\mathbf{R}^1(f_{\mathrm{DR}*}(\mathrm{Ad}(P)))$$

which may be identified with the tangent bundle of $\mathcal{Y}_{X/S}^{\rho}$ over S at the S -valued point (P, ∇_P) . Now, by considering what it means to construct a deformation of (P, ∇_P) as a torally crys-stable bundle of level l by gluing together local deformations, it is clear that the tangent bundle of $\mathcal{Y}_{X/S}^{\rho;l}$ over S at (P, ∇_P) is given by the image of the first cohomology module of the subcomplex

$$F^0(\mathrm{Ad}(P)) \rightarrow \mathrm{Ad}^q(P) \otimes \omega_{X^{\log}/S^{\log}}$$

in the first cohomology module of the complex $\mathrm{Ad}(P) \rightarrow \mathrm{Ad}^q(P) \otimes \omega_{X^{\log}/S^{\log}}$, i.e.,

$$F^0(\mathbf{R}^1(f_{\mathrm{DR}*}(\mathrm{Ad}(P))))$$

Moreover, the deformations for which the Kodaira-Spencer locus remains fixed clearly correspond to deformations arising from the first cohomology module of the complex

$$F^1(\mathrm{Ad}(P)) \rightarrow F^0(\mathrm{Ad}^q(P)) \otimes \omega_{X^{\log}/S^{\log}}$$

i.e., $F^1(\mathbf{R}^1(f_{\mathrm{DR}})_*(\mathrm{Ad}(P)))$.

Thus, it remains only to verify that any deformation of the Kodaira-Spencer locus may be realized by some element of $F^0(\mathbf{R}^1(f_{\mathrm{DR}})_*(\mathrm{Ad}(P)))$. But the obstruction class to constructing such an element forms an element of the second cohomology module of the complex

$$F^1(\mathrm{Ad}(P)) \rightarrow F^0(\mathrm{Ad}^q(P)) \otimes \omega_{X^{\log}/S^{\log}}$$

which is the cokernel of the natural morphism

$$\mathbf{R}^1 f_*(F^1(\mathrm{Ad}(P))) = \mathbf{R}^1 f_*(\mathcal{L}) \rightarrow \mathbf{R}^1 f_*(F^0(\mathrm{Ad}^q(P)) \otimes \omega_{X^{\log}/S^{\log}})$$

induced by the connection on $\mathrm{Ad}(P)$. Here, $\mathcal{L} \stackrel{\mathrm{def}}{=} h^* \omega_{P/X} \cong \omega_{X^{\log}/S^{\log}}(-D_\kappa)$ – cf. Proposition 1.7. Moreover, we have an exact sequence

$$0 \rightarrow \mathcal{L}(-M_f) = F^1(\mathrm{Ad}^q(P)) \rightarrow F^0(\mathrm{Ad}^q(P)) \rightarrow \mathcal{O}_X = (F^0/F^1)(\mathrm{Ad}^q(P)) \rightarrow 0$$

Next, observe that $\mathbf{R}^1 f_*(\mathcal{L}(-M_f) \otimes \omega_{X^{\log}/S^{\log}})$ is dual to $f_* \mathcal{L}^{-1}$, which is zero since \mathcal{L} is *ample*. Moreover, $\mathbf{R}^1 f_*(F^1(\mathrm{Ad}(P))) = \mathbf{R}^1 f_*(\omega_{X^{\log}/S^{\log}}(-D_\kappa))$ surjects onto $\mathbf{R}^1 f_*(\omega_{X^{\log}/S^{\log}})$ (since $X \rightarrow S$ is of relative dimension 1). This shows that the second cohomology module in question is identically zero, hence that the obstruction class under consideration is itself zero. This observation completes the proof of Lemma 3.8. \square

Before proceeding, we would like to discuss a certain technical result concerning specializations of sections of \mathbf{P}^1 -bundles which we will use in Chapter II. Let S be the spectrum of a discrete valuation ring A . For simplicity, we assume that the residue field of A is algebraically closed. Let $f^{\log} : X^{\log} \rightarrow S^{\log}$ be an r -pointed stable log-curve of genus g over S which is generically smooth. Let $(\pi : P \rightarrow X, \nabla_P)$ be a \mathbf{P}^1 -bundle with connection over X . Let us assume that over the generic point η of S , the bundle $(P_\eta, \nabla_{P_\eta})$ is *totally crys-stable of level $l > 0$ on X_η^{\log}* , with Hodge section $h_\eta : X_\eta \rightarrow P_\eta$. Let $D_\eta \subseteq X_\eta$ be the Kodaira-Spencer locus of $(P_\eta, \nabla_{P_\eta})$.

Proposition 3.9. *Suppose that $D_\eta \subseteq X_\eta$ extends to an S -flat, l -balanced divisor $D \subseteq X$. Let $D_s \stackrel{\mathrm{def}}{=} X_s \cap D$ be the restriction of D to the special fiber X_s of X .*

Let us also assume that $h_\eta : X_\eta \rightarrow P_\eta$ extends to a section $h_U : U \rightarrow P$ (where $U \stackrel{\text{def}}{=} X - D_s$) whose associated Kodaira-Spencer morphism is an isomorphism over $U - D_\eta$. Then (P, ∇_P) is totally crys-stable on X^{\log} , and its restriction to X_s is of level $\geq l$. Moreover, the Hodge section $h_s : X_s \rightarrow P_s$ of $(P, \nabla_P)|_{X_s}$ extends $h_U|_{X_s}$.

Proof. First, by passing to a log-admissible (in the sense of [Mzk2], §3) covering of X^{\log} of degree a power of 2, one sees easily that it suffices to prove the Proposition in the case where there exists a line bundle \mathcal{F} on P whose square is $\tau_{P/X}$. Let $(\mathcal{E}, \nabla_{\mathcal{E}})$ be the rank two vector bundle with connection given by $\pi_*\mathcal{F}$, together with the connection induced by ∇_P . Note that $(\mathcal{E}, \nabla_{\mathcal{E}})$ has trivial determinant. Let $\mathcal{L}_U = h_U^*\mathcal{F}^{-1}$. Since X is regular at D_s , \mathcal{L}_U extends uniquely to a line bundle \mathcal{L} on X . Since $\mathcal{L}_U^{\otimes 2} \cong \omega_{X^{\log}/S^{\log}}(-D)|_U$ (by the assumption concerning the Kodaira-Spencer morphism of h_U), it follows that $\mathcal{L}^{\otimes 2} \cong \omega_{X^{\log}/S^{\log}}(-D)$. Now observe that h_U defines a locally split inclusion $\mathcal{L}_U \hookrightarrow \mathcal{E}_U$. Since X is normal of dimension two at D_s , this inclusion extends to an inclusion $\iota : \mathcal{L} \hookrightarrow \mathcal{E}$, which, however, may not be locally split over all of X . Nevertheless, we can still form the composite of $\iota : \mathcal{L} \hookrightarrow \mathcal{E}$ with $\mathcal{E} \rightarrow \mathcal{E} \otimes_{\mathcal{O}_X} \omega_{X^{\log}/S^{\log}} \rightarrow \mathcal{L}^{-1} \otimes_{\mathcal{O}_X} \omega_{X^{\log}/S^{\log}}$ (where the second morphism is $\nabla_{\mathcal{E}}$ composed with the dual of ι). This composite $\mathcal{L} \rightarrow \mathcal{L}^{-1} \otimes_{\mathcal{O}_X} \omega_{X^{\log}/S^{\log}}$ is clearly linear and vanishes on D_η , hence on D . Thus, we obtain a morphism of line bundles $\phi : \mathcal{L} \rightarrow \mathcal{L}^{-1} \otimes_{\mathcal{O}_X} \omega_{X^{\log}/S^{\log}}(-D)$ which (by assumption) is an isomorphism outside D_s . Since X is normal at D_s , it follows that ϕ is an isomorphism. Now let us suppose that the inclusion ι_s factors through a *locally split* inclusion $\mathcal{M} \hookrightarrow \mathcal{E}_s$. Then it is clear that the zero locus of the Kodaira-Spencer morphism $\mathcal{M}^{\otimes 2} \rightarrow \omega_{X^{\log}/S^{\log}}$ is contained in the zero locus of the composite of $\mathcal{L}_s^{\otimes 2} \cong (\omega_{X^{\log}/S^{\log}}(-D))|_{X_s}$ (formed from ϕ_s) with the inclusion $\omega_{X^{\log}/S^{\log}}(-D) \hookrightarrow \omega_{X^{\log}/S^{\log}}$ (restricted to X_s). Thus, the zero locus of the Kodaira-Spencer morphism of $\mathcal{M} \hookrightarrow \mathcal{E}_s$ is contained in D_s , hence is balanced (since D_s is balanced). It thus follows that the morphism $h_s : X_s \rightarrow P_s$ defined by $\mathcal{M} \hookrightarrow \mathcal{E}_s$ has canonical height $\leq -l$. Moreover, the restriction of h_s to every irreducible component of X_s has *negative canonical height*. It thus follows immediately that (P, ∇_P) is totally crys-stable. \circ

Let l be any nonnegative half-integer. In the universal case, we let $\overline{\mathcal{Y}}_{g,r}^{\rho;l} \rightarrow \overline{\mathcal{M}}_{g,r}$ be the relative algebraic space that we have been calling $\mathcal{Y}_{X/S}^{\rho;l} \rightarrow S$. We summarize the results of this subsection as follows:

Theorem 3.10. *The algebraic stack $\overline{\mathcal{Y}}_{g,r}^{\rho;l}$ is smooth of relative dimension $2(3g - 3 + r) - 2l - (g - 1)$ (respectively, $2(3g - 3 + r)$) over $\overline{\mathcal{M}}_{g,r}$ if $0 < l \leq \chi$ (respectively, $l = 0$), and empty if $l > \chi$. Moreover, the natural morphism*

$$\coprod_{l \geq 0} \overline{\mathcal{Y}}_{g,r}^{\rho;l} \rightarrow \overline{\mathcal{Y}}_{g,r}^{\rho}$$

is a monomorphism such that the restrictions $\overline{\mathcal{Y}}_{g,r}^{\rho;l} \rightarrow \overline{\mathcal{Y}}_{g,r}^{\rho}$ to the components on the left are immersions with the following property: the closure of the image of $\overline{\mathcal{Y}}_{g,r}^{\rho;l}$ in $\overline{\mathcal{Y}}_{g,r}^{\rho}$ is disjoint from $\overline{\mathcal{Y}}_{g,r}^{\rho;l'}$ for $l' < l$. Finally, if $l = \chi$, then the immersion $\overline{\mathcal{Y}}_{g,r}^{\rho;l} \rightarrow \overline{\mathcal{Y}}_{g,r}^{\rho}$ is a closed immersion.

We shall refer to the collection of natural morphisms appearing in the preceding paragraph as the Hodge structure on $\overline{\mathcal{Y}}_{g,r}^{\rho} \rightarrow \overline{\mathcal{M}}_{g,r}$.

Proof. It remains only to prove that the assertions concerning immersions. Let us first prove that the morphisms $\overline{\mathcal{Y}}_{g,r}^{\rho;l} \rightarrow \overline{\mathcal{Y}}_{g,r}^{\rho}$ are immersions. We may assume that $l > 0$. By Corollary 2.13, it suffices to check that the valuative criterion for radimmersions is satisfied. Thus, let S^{\log} be $\text{Spec}(A)$, where A is a discrete valuation ring, equipped with the log structure defined by the special fiber. Let $X^{\log} \rightarrow S^{\log}$ be an r -pointed stable log-curve of genus g which is generically smooth. (Note that it suffices to consider the generically smooth case since in the arbitrary case, we can always break up a stable $X^{\log} \rightarrow S^{\log}$ into generically smooth components.) Let (P, ∇_P) be a torally crys-stable bundle on X^{\log} whose restrictions to the generic and special fibers are of level l . We would like to show that (P, ∇_P) itself is of level l . Let $h_{\eta} : X_{\eta} \rightarrow P_{\eta}$ be a Hodge section of canonical height $-l$. Let $Y \subseteq P$ be the schematic closure of the image of h_{η} . Then Y is S -flat, and the composite $Y \rightarrow P \rightarrow X$ is a proper morphism $\phi : Y \rightarrow X$ which is an isomorphism over η . Moreover, $h_{\eta} \circ \phi_{\eta}$ extends to a morphism $h_Y : Y \rightarrow P$. Note that the special fiber Y_s has two types of irreducible components: irreducible components that map isomorphically to irreducible components of X_s , which we shall call *strict components*, and irreducible components that map to a point of X_s , which we shall call *exceptional components*. Let $\mathcal{L} = h_Y^* \tau_{P/X}$. Since $\deg(\mathcal{L}_{\eta}) = -2l$ and $Y \rightarrow S$ is flat, it follows that $\deg(\mathcal{L}_s) = -2l$. Note, moreover, that the restriction of \mathcal{L}_s to an exceptional component clearly has positive degree.

On the other hand, since $(P, \nabla_P)_s$ is crys-stable of level l , it has a Hodge section $h' : X_s \rightarrow P_s$ of canonical height $-l$. Let $\mathcal{M} = (h')^* \tau_{P/X}$. Now we *claim* that for any irreducible component $Z \subseteq Y_s$, $\deg(\mathcal{L}_s|_Z) \geq \deg(\mathcal{M}|_Z)$. Indeed, since the canonical height of the restriction of h' to any irreducible component of Y_s is ≤ 0 , it follows that it suffices to check this inequality for those Z such that $\deg(\mathcal{L}_s|_Z) < 0$. Let $Z \subseteq Y_s$ be any such component (which is necessarily strict). Since a \mathbf{P}^1 -bundle over an irreducible curve can have at most one section of negative canonical height, it follows that the restriction of h_Y to Z is equal to

the restriction of h' to Z . Thus, for such Z , the claimed inequality is, in fact, an equality. This completes the proof of the claim.

Thus, in summary, summing over all $Z \subseteq Y_s$, we obtain that $-2l = \deg(\mathcal{L}_s) \geq \deg(\mathcal{M}) = -2l$. Thus, this inequality must be an equality. In particular, there cannot exist $Z \subseteq Y_s$ for which $\deg(\mathcal{L}_s|_Z) > 0$. But this means that there cannot exist exceptional components, i.e., $Y = X$, so h_η extends to $h : X \rightarrow P$ whose restriction to X_s is h' . This completes the proof of the verification of the valuative criterion for radimmersions.

The assertion concerning the disjointness of the closure of the image of $\overline{\mathcal{Y}}_{g,r}^{\rho;l}$ in $\overline{\mathcal{Y}}_{g,r}^\rho$ is from $\overline{\mathcal{Y}}^{\rho;l'}$ for $l' < l$ is proven as follows: If $l' > 0$, one simply mimicks the preceding valuative proof, except that this time we assume that h' has canonical height $-l'$; as above, one obtains the inequality $-2l \geq -2l'$, i.e., $l' \geq l$, which is absurd. If $l' = 0$, without even introducing h' , one obtains the inequality $-2l \geq 0$, which is again absurd. This completes the proof of this assertion.

Finally, let us consider the case $l = \chi$. In this case, it suffices to verify the valuative criterion for properness. The proof is similar to the preceding proof involving the valuative criterion for radimmersions, except that we cannot assume that we are given a morphism h' . Thus, we maintain the notations of the preceding proof. For every irreducible component $Z \subseteq Y_s$, let $\chi_Z \stackrel{\text{def}}{=} \frac{1}{2} \deg(\omega_{Z/\log/s/\log})$ if Z is strict (hence can be regarded as an irreducible component of X_s ; the log structure is then given by regarding as marked points on Z the points of Z that map to nodes or marked points of X_s); let $\chi_Z = 0$ if Z is exceptional. This time, let us verify (using the assumption $l = \chi$) that for every $Z \subseteq Y_s$, we have $\deg(\mathcal{L}_s|_Z) \geq -2\chi_Z$. Just as above, it suffices to check this inequality for those Z such that $\deg(\mathcal{L}_s|_Z) < 0$. Let $Z \subseteq Y_s$ be any such component (which is necessarily strict). On the other hand, if $\deg(\mathcal{L}_s|_Z) < -2\chi_Z$, then (by degree considerations) its Kodaira-Spencer morphism must vanish, thus contradicting the fact that $(P, \nabla_P)_s$ is crys-stable. This proves the inequality $\deg(\mathcal{L}_s|_Z) \geq -2\chi_Z$. Summing over Z , we thus obtain $-2\chi = \deg(\mathcal{L}_s) \geq -2\chi$ (since the sum of the χ_Z 's over all Z is χ). Just as before, this implies that exceptional components cannot exist. This completes the verification of the the valuative criterion for properness, and hence the proof of the Theorem. \circ

§3.3. Clutching Behavior

Let n be a natural number (≥ 1). Suppose that we are given the following data:

- (1) a finite graph Γ of n vertices, numbered 1 through n ;

- (2) for each $i = 1, \dots, n$, a pair of nonnegative integers (g_i, r_i) such that $2g_i - 2 + r_i \geq 1$;
- (3) for each $i = 1, \dots, n$, an injection $\lambda_i : E_i \hookrightarrow \{1, \dots, r_i\}$, where E_i is the set of ends of edges emanating from the i^{th} vertex of the graph Γ . (Thus, if an edge runs from the i^{th} vertex to itself, it defines *two* elements of E_i ; if an edge runs from the i^{th} vertex to a different vertex, it defines *one* element of E_i .)

Once this *combinatorial data* has been specified, we can use it to glue together any collection

$$\{f_i^{\log} : X_i^{\log} \rightarrow S^{\log}\}_{i=1, \dots, n}$$

of curves (where the curve numbered i is an r_i -pointed stable curve of genus g_i) to obtain a new pointed, stable curve $f^{\log} : X^{\log} \rightarrow S^{\log}$ in such a way that the *dual graph* of f^{\log} is given by Γ , that is:

- (1) the vertex i corresponds to $f_i^{\log} : X_i^{\log} \rightarrow S^{\log}$, an irreducible component of X^{\log} ;
- (2) if ϵ is an edge with ends ϵ_1 (attached to vertex i) and ϵ_2 (attached to vertex j) such that $\lambda_i(\epsilon_1) = a$ and $\lambda_j(\epsilon_2) = b$, then ϵ corresponds to a node on X^{\log} obtained by gluing together X_i^{\log} at the a^{th} marked point to X_j^{\log} at the b^{th} marked point;
- (3) we order the marked points of $f^{\log} : X^{\log} \rightarrow S^{\log}$ *lexicographically*: i.e., a marked point lying on X_i^{\log} comes before a marked point lying on X_j^{\log} (if $i < j$), and among marked points lying on X_i^{\log} , we take the ordering induced by the original ordering of marked points on X_i^{\log} .

The genus, g , and the number of marked points, r , of $f^{\log} : X^{\log} \rightarrow S^{\log}$ can then be computed combinatorially solely from Γ , the g_i 's, the r_i 's and the λ_i 's. Let us denote by $\mu_i^{\log} : X_i^{\log} \rightarrow X^{\log}$ the natural morphism obtained from the construction of X^{\log} by gluing together the X_i^{\log} 's.

Definition 3.11. We shall call the combinatorial data

$$\mathcal{D} \stackrel{\text{def}}{=} \{\Gamma; \{(g_i, r_i)\}_{i=1, \dots, r}; \{\lambda_i\}_{i=1, \dots, r}\}$$

clutching data for an r -pointed curve of genus g .

By carrying out the above construction in the universal case, we obtain a morphism of moduli stacks:

$$\kappa_{\mathcal{D}} : \prod_{i=1}^n \overline{\mathcal{M}}_{g_i, r_i} \rightarrow \overline{\mathcal{M}}_{g, r}$$

which we call the *clutching morphism associated to \mathcal{D}* .

Now let us suppose that we are given a crys-stable bundle $(\pi : P \rightarrow X; \nabla_P)$ on X^{\log} . Then it is not necessarily the case that $(\mu_i^{\log})^*(\pi : P \rightarrow X; \nabla_P)$ will be crys-stable on X_i^{\log} . The problem is that since in general, marked points of X_i^{\log} might be sent to nodes of X^{\log} (and not to marked points), there is no reason why the monodromy at such marked points of X_i^{\log} should be nilpotent. We therefore make the following

Definition 3.12. If the $(\mu_i^{\log})^*(\pi : P \rightarrow X; \nabla_P)$ are crys-stable on X_i^{\log} for all i , then we say that $(\pi : P \rightarrow X; \nabla_P)$ is of *restrictable type*.

Now let us suppose that we are given torally crys-stable bundles $(\pi_i : P_i \rightarrow X_i; \nabla_{P_i})$ on X_i^{\log} , together with, for each edge ϵ of the graph Γ (with ends ϵ_1 and ϵ_2 , attached to vertices i and j , respectively) an isomorphism of \mathbf{P}^1 -bundles $\xi_{\epsilon} : \sigma_{i, \lambda_i(\epsilon_1)}^* P_i \rightarrow \sigma_{j, \lambda_j(\epsilon_2)}^* P_j$ over S that respects the monodromy operators operating on either side. Then we can glue together the $(P_i; \nabla_{P_i})$'s by means of the ξ_{ϵ} 's at the marked points to obtain a new torally crys-stable bundle $(\pi : P \rightarrow X; \nabla_P)$ on X^{\log} . By construction, if each of the $(P_i; \nabla_{P_i})$'s is crys-stable, then $(P; \nabla_P)$ is crys-stable of *restrictable type*, and in fact, it is easy to see that every crys-stable bundle on X^{\log} of restrictable type is obtained in this way.

Moreover, this gluing procedure respects the various Hodge morphisms of Theorem 3.10 in the following sense: If all of the (P_i, ∇_{P_i}) 's define points of $\mathcal{Y}^{\rho; 0}$, then so will (P, ∇_P) . If (P_i, ∇_{P_i}) defines a point of $\mathcal{Y}^{\rho; N_i}$ (where N_i is a *positive* half-integer) for all i , and the ξ_{ϵ} 's respect the sections of negative canonical height on either side, then (P, ∇_P) will define a point of $\mathcal{Y}^{\rho; N}$, where $N = \sum_i N_i$.

Of course, it is possible to translate the last two paragraphs into the language of morphisms among various \mathcal{Y}^{ρ} 's and $\mathcal{Y}^{\rho; N}$'s, but we shall leave this to the reader.

Example 3.13. Finally, it is worth taking a look at what happens in the *totally degenerate case*, i.e., the case when all the g_i are zero, and all the r_i are 3. Let (P, ∇_P) be a torally crys-stable bundle on X^{\log} of radii ρ . Let (P_i, ∇_{P_i}) be its restriction to X_i^{\log} . Then the \mathbf{P}^1 -bundle $P_i \rightarrow X_i$ is either trivial, or of the form $\mathbf{P}(\mathcal{O}_{X_i} \oplus \omega_{X^{\log}/S^{\log}})$. This separates the space

of torally crys-stable bundles of radii ρ on X^{\log} into precisely $2^{(2g-2+r)}$ connected components (since $n = 2g - 2 + r$). The space of connections on $P_i \rightarrow X_i$ (up to isomorphism) with nonzero monodromy at the marked points is easily seen to be smooth and $(3 - \zeta_i)$ -dimensional, where ζ_i is the number of marked points of X_i^{\log} that are also marked points of X^{\log} . Thus, once one fixes ∇_{P_i} , in order to reconstruct (P, ∇_P) , it suffices to give the clutching isomorphisms $\xi_\epsilon : \sigma_{i, \lambda_i(\epsilon)}^* P_i \rightarrow \sigma_{j, \lambda_j(\epsilon)}^* P_j$ that respect the monodromy on both sides. It is easy to see that, for each edge ϵ , the space of such ξ_ϵ is smooth of dimension 1. Thus, tallying everything up, we see that the relative dimension of $\mathcal{Y}_{X/S}^\rho$ over S is given by twice the number of nodes of X^{\log} , i.e., $2(3g - 3 + r)$. This agrees with Theorem 2.7. Suppose now that all the ρ_i are zero. Then the relative dimension over S of the subspace of crys-stable bundles on X^{\log} of *restrictable type* is given by the number of nodes of X^{\log} , i.e., $3g - 3 + r$ (since such a bundle is completely determined by the ξ_ϵ 's).

§4. Torally Indigenous Bundles

In this §, we define torally indigenous bundles, and discuss their basic properties. Since many of the proofs are similar to the “classical case” (discussed in great detail in [Mzk1], Chapter I), we shall often refer to [Mzk1], Chapter I for proofs. In what follows, we shall call the indigenous bundles dealt with in [Mzk1], Chapter I, *classical indigenous bundles* to distinguish them from “torally indigenous bundles.” As before, we let S^{\log} be a fine noetherian log scheme on which p (an odd prime) is nilpotent. Let $f^{\log} : X^{\log} \rightarrow S^{\log}$ be an r -pointed stable curve of genus g (where $2g - 2 + r \geq 1$). Let $\sigma_1, \dots, \sigma_r : S \rightarrow X$ be the r marked points, and let $D \subseteq X$ be the divisor of marked points.

§4.1. Definitions

As before, we fix for each $i = 1, \dots, r$, a *radius* $\rho_i \in \Gamma(S, \mathcal{O}_S)$, which is either 0 or a section of \mathcal{O}_S^\times . Let $\pi : P \rightarrow X$ be a \mathbf{P}^1 -bundle on X , and let ∇_P be a logarithmic connection on this \mathbf{P}^1 -bundle.

Definition 4.1. We shall call $(\pi : P \rightarrow X, \nabla_P)$ a *torally indigenous bundle* (of radii $\{\rho_i\}$) if there exists a section $\sigma_H : X \rightarrow P$ at which the Kodaira-Spencer morphism is an isomorphism, and the monodromy operator μ_i satisfies $\langle \mu_i, \mu_i \rangle = 2\rho_i^2$, for all $i = 1, \dots, r$.

Remark. In other words, a torally indigenous bundle is the same thing as a torally crys-stable bundle of level $\chi = \frac{1}{2}(2g - 2 + r)$ (cf. Lemma 3.4).

Remark. Note that when a radius ρ_i is zero, the definition implies that the connection has nilpotent monodromy at σ_i . Thus, when all the ρ_i are zero, one recovers the notion of a *classical indigenous bundle* (i.e., what we called an “indigenous bundle” in [Mzk1], Chapter I, Definition 2.2).

By the same proof as in the classical case, one sees that σ_H , if it exists, is unique, and of canonical height $1 - g - \frac{1}{2}r$. We shall refer to it as the *Hodge section*. Moreover, a torally indigenous bundle defines a *crystalline Schwarz structure* (as in [Mzk1], Chapter I, Definition 1.5). Hence, by [Mzk1], Chapter I, Proposition 1.4, we obtain (in the notation of *loc. cit.*) a natural isomorphism

$$P \cong \mathbf{P}(\mathcal{J}/\mathcal{J}^{[3]})$$

Although we will not review all of the details concerning the construction of this isomorphism here, because of its importance, we will give a brief “conceptual” review of this construction. Let $\mathcal{O}_{\widehat{\mathcal{X}}^{\text{bi}}}$ be the sheaf of *biformal functions* on X . This sheaf is constructed by forming the completion of the PD-envelope of the diagonal in $X^{\log} \times_{S^{\log}} X^{\log}$. (See [Mzk1]^{*}, Chapter I, §1, for more details.) There are two injections of \mathcal{O}_X into $\mathcal{O}_{\widehat{\mathcal{X}}^{\text{bi}}}$, one “from the left” and one “from the right.” Following [Mzk1], we regard the image of the \mathcal{O}_X on the right as “constants” (in the terminology of [Mzk1], Chapter I, *biformal constants*), i.e., the result of base-changing via $X \rightarrow S$. Thus, the image of the \mathcal{O}_X on the left consists of “true functions.” By integrating (P, ∇_P) , one obtains a subsheaf $\mathcal{S} \subseteq \mathcal{O}_{\widehat{\mathcal{X}}^{\text{bi}}}$ (valued in the category of *sets*) with the property that (roughly speaking) any two functions in this subsheaf are related by a linear fractional transformation whose coefficients are biformal constants. Such a subsheaf is called a *Schwarz structure*. Moreover, the theory of [Mzk1], Chapter I, allows one to reconstruct (P, ∇_P) from the subsheaf \mathcal{S} . Roughly speaking, the point (cf. [Mzk1], Chapter I, Proposition 1.4, for more details) is that the \mathbf{P}^1 -bundle $P \rightarrow X$ may be reconstructed as precisely the \mathbf{P}^1 -bundle on which the sections of \mathcal{S} are the degree one rational functions. Moreover, the surjection $\mathcal{O}_{\widehat{\mathcal{X}}^{\text{bi}}} \rightarrow \mathcal{O}_X$ (arising from the diagonal embedding) defines a special “point” of this \mathbf{P}^1 -bundle which, in fact, turns out to be the Hodge section. Finally, one observes that the global differentials on this bundle of \mathbf{P}^1 ’s (i.e., relative differentials for $P \rightarrow X$) that are regular everywhere, except (possibly) for a pole of order ≤ 3 at this special point forms a vector bundle of rank two on X whose projectivization can be canonically identified with $P \rightarrow X$. On the other hand, this vector bundle is naturally dual to the bundle of “Taylor expansions to order two of functions vanishing at the special point” (i.e., $\mathcal{J}/\mathcal{J}^{[3]}$, where \mathcal{J} is the kernel of the surjection $\mathcal{O}_{\widehat{\mathcal{X}}^{\text{bi}}} \rightarrow \mathcal{O}_X$). Indeed, the duality pairing is given by multiplying such a differential by such a Taylor expansion

and then taking the residue at the special point. Putting these last two facts together, we see that we get a canonical isomorphism of $P \rightarrow X$ with the projectivization of $\mathcal{J}/\mathcal{J}^{[3]}$.

§4.2. Explicit Computation of Monodromy

Now, let us see what the canonical toral structure at a marked point (cf. the paragraphs preceding Proposition 1.6) corresponds to under the above isomorphism $P \cong \mathbf{P}(\mathcal{J}/\mathcal{J}^{[3]})$. Let ϕ be a relative rational function for $\sigma_i^*P \rightarrow S$ which has a simple zero at $q_i[0]$, and a simple pole at $q_i[\infty]$ (and no other zeroes or poles). By passing from relative rational functions to sections of the Schwarz structure associated to our torally indigenous bundle, we see that such a ϕ corresponds to a section $z \in \mathcal{O}_{\widehat{\mathcal{X}}^{\text{bi}}} \widehat{\otimes}_{\mathcal{O}_{X, \sigma_i^{-1}}} \mathcal{O}_S$. (Here, $\mathcal{O}_{\widehat{\mathcal{X}}^{\text{bi}}}$ is the sheaf of *biformal functions* in the terminology of [Mzk1], Chapter I, §1.) Let us consider the monodromy action on z : if we apply the monodromy operator (for biformal functions) M to z , we get a section $M(z)$ of $\mathcal{O}_{\widehat{\mathcal{X}}^{\text{bi}}} \widehat{\otimes}_{\mathcal{O}_{X, \sigma_i^{-1}}} \mathcal{O}_S$. Since the correspondence between ϕ and z preserves the monodromy action, we compute that

$$M(z) = 2\rho_i \cdot z$$

But this determines z , up to multiplication by a section of \mathcal{O}_S^\times . Indeed, let t be a local section of \mathcal{O}_X that defines the subscheme which is the image of σ_i , and let

$$\delta = 1 - \left(\frac{1 \otimes t}{t \otimes 1} \right) \in \mathcal{O}_{\widehat{\mathcal{X}}^{\text{bi}}} \widehat{\otimes}_{\mathcal{O}_{X, \sigma_i^{-1}}} \mathcal{O}_S$$

Then it is immediate (from the fact that M is a derivation satisfying $M(1 - \delta) = \delta - 1$) that z must be an invertible multiple of

$$\left(\frac{1 \otimes t}{t \otimes 1} \right)^{-2\rho_i} \stackrel{\text{def}}{=} \exp\{-2\rho_i \cdot \log(1 - \delta)\}$$

In the following Definition, let us assume that z is actually equal to the expression above.

Definition 4.2. When ρ_i is nonzero, we shall call that endomorphism of $\mathbf{P}(\mathcal{J}/\mathcal{J}^{[3]})|_{\sigma_i}$ given by taking the image of $z - 1$ (respectively, $z^{-1} - 1$) in $\mathcal{J}/\mathcal{J}^{[3]}$ to ρ_i (respectively, $-\rho_i$) times itself, the *standard monodromy endomorphism of radius ρ_i* . When $\rho_i = 0$, we shall call that endomorphism of $\mathbf{P}(\mathcal{J}/\mathcal{J}^{[3]})|_{\sigma_i}$ given by taking the image of $\{\log(1 - \delta)\}^2$ (respectively, $\log(1 - \delta)$) in $\mathcal{J}/\mathcal{J}^{[3]}$ to the image of $-\log(1 - \delta)$ (respectively, 0) in $\mathcal{J}/\mathcal{J}^{[3]}$, the *standard monodromy endomorphism of radius zero*.

Remark. The reader might find it strange that, in the above Definition, in the case where ρ_i is nonzero, we took the standard monodromy endomorphism to be multiplication by $\pm\rho_i$ rather than $\pm 2\rho_i$. The reason for this apparent discrepancy is that the way that we constructed (in [Mzk1], Chapter I, §1) a connection on $\mathbf{P}(\mathcal{J}/\mathcal{J}^{[3]})$ out of a crystalline Schwarz structure was *not* by taking a section ϕ of the Schwarz structure, applying the “biformal function connection” to ϕ , and then looking at the image of ϕ in $\mathcal{J}/\mathcal{J}^{[3]}$; indeed, it is easy to see that this sort of definition is not well-defined. Instead, the approach we took was somewhat more indirect (see *loc. cit.* for more details). Thus, if one traces through the correct definition of the connection on $\mathbf{P}(\mathcal{J}/\mathcal{J}^{[3]})$, one sees that multiplication by $\pm\rho_i$ (*not* $\pm 2\rho_i$) is indeed what we want.

Next, we would like to observe that, relative to the canonical isomorphism $P \cong \mathbf{P}(\mathcal{J}/\mathcal{J}^{[3]})$ discussed at the end of the preceding subsection, the monodromy endomorphism of (P, ∇_P) corresponds precisely to the “standard monodromy endomorphism” of Definition 4.2. We begin with the case of nonzero radius. First, observe that the image of the Schwarz structure associated to (P, ∇_P) in $\mathcal{O}_{\widehat{\mathcal{X}}^{\text{bi}}} \widehat{\otimes}_{\mathcal{O}_{X, \sigma_i^{-1}}} \mathcal{O}_S$ is generated by z . Thus, one may think of the \mathbf{P}^1 -bundle $P \rightarrow X$ as precisely the \mathbf{P}^1 -bundle on which z serves as a “standard coordinate.” Relative to this coordinate, the Hodge section is the point $z = 1$. The differentials with poles of order ≤ 3 at this point are given by

$$(a + bz) \cdot \frac{dz}{(z-1)^3}$$

where $a, b \in \mathcal{O}_S$. Relative to the duality between differentials and Taylor expansions, the point $z = \infty$ (which is the zero locus of the differential $\frac{dz}{(z-1)^3}$) corresponds to the “line” of Taylor expansions generated by $z-1$ (since $(z-1) \cdot \frac{dz}{(z-1)^3}$ has residue zero at $z = 1$). Similarly, the point $z = 0$ (which is the zero locus of the differential $\frac{z dz}{(z-1)^3}$) corresponds to the “line” of Taylor expansions generated by $z^{-1} - 1$ (since $(z^{-1} - 1) \cdot \frac{z dz}{(z-1)^3}$ has residue zero at $z = 1$). Since (by definition) the eigenspaces of the monodromy operator on P are the subspaces corresponding to the unique zero and the unique pole of z , and we know what the respective eigenvalues are, this completes the proof of the following result in the toral case:

Proposition 4.3. *Relative to the above isomorphism $P \cong \mathbf{P}(\mathcal{J}/\mathcal{J}^{[3]})$, the monodromy endomorphism of $\mathbf{P}(\mathcal{J}/\mathcal{J}^{[3]})|_{\sigma_i}$ arising from ∇_P is the standard one of radius ρ_i . In particular, if ρ_i is nonzero, then the section $q_i[0]$ (respectively, $q_i[\infty]$) corresponds to the rank one \mathcal{O}_S -submodule of $\mathcal{J}/\mathcal{J}^{[3]}|_{\sigma_i}$ generated by the image of $z-1$ (respectively, $z^{-1}-1$) in $\mathcal{J}/\mathcal{J}^{[3]}|_{\sigma_i}$.*

Proof. It remains to consider the case of radius 0. The argument is entirely similar to the toral case. Indeed, let $z \stackrel{\text{def}}{=} \log(1 - \delta)$. Then one may think of the \mathbf{P}^1 -bundle $P \rightarrow X$ as precisely the \mathbf{P}^1 -bundle on which z serves as a “standard coordinate.” Relative to this coordinate, the Hodge section is the point $z = 0$. The differentials with poles of order ≤ 3 at this point are given by

$$(a + bz) \cdot \frac{dz}{z^3}$$

where $a, b \in \mathcal{O}_S$. Relative to the duality between differentials and Taylor expansions, the point $z = \infty$ (which is the zero locus of the differential $\frac{dz}{z^3}$) corresponds to the “line” of Taylor expansions generated by z (since $z \cdot \frac{dz}{z^3}$ has residue zero at $z = 0$). Similarly, the point $z = a$ (which is the zero locus of the differential $\frac{(z-a) dz}{z^3}$) corresponds to the “line” of Taylor expansions generated by $az + z^2$ (since $(az + z^2) \cdot \frac{(z-a) dz}{z^3}$ has residue zero at $z = 0$). Since we know that the monodromy operator on P is nilpotent with respect to the subspace corresponding to the point $z = \infty$ and maps a to $a - 1$, the result follows immediately. \square

§4.3. Moduli and de Rham Cohomology

It thus follows, by dealing with *normalized connections* on $\mathbf{P}(\mathcal{I}/\mathcal{I}^{[3]})$ (as in [Mzk1], Chapter I, Definition 1.10), that the obstruction to the existence of a torally indigenous bundle with the given radii is a torsor over $(\omega_{X^{\log}/S^{\log}})^{\otimes 2}(-D)$ (where $D \subseteq X$ is the divisor of marked points). Thus, just as in the classical case (cf. the discussion preceding [Mzk1], Chapter I, Theorem 1.9), we have the following result:

Theorem 4.4. *Étale locally on S , there always exists a torally indigenous bundle of radii $\{\rho_i\}$ on X^{\log} . Moreover, the functor that assigns to $T^{\log} \rightarrow S^{\log}$ the set of isomorphism classes of torally indigenous bundles of radii $\{\rho_i\}$ on $X_T^{\log} \stackrel{\text{def}}{=} X^{\log} \times_{S^{\log}} T^{\log}$ is a torsor over $f_*(\omega_{X^{\log}/S^{\log}})^{\otimes 2}(-D)$.*

Let $\overline{\mathcal{M}}_{g,r}$ denote the moduli stack of r -pointed stable curves of genus g over \mathbf{Z}_p . Let $\rho = \{\rho_i\}$, and suppose (just for this paragraph) that all the ρ_i ’s are in \mathbf{Z}_p . Then, by working in the universal case, we get a torsor, called the *Schwarz torsor*, over $\Omega_{\overline{\mathcal{M}}_{g,r}}^{\log}$

$$\overline{\mathcal{S}}_{g,r}^{\rho} \rightarrow \overline{\mathcal{M}}_{g,r}$$

parametrizing the isomorphism classes of torally indigenous bundles with radii given by ρ on the universal curve. Finally, just as in the classical case, we have the following:

Proposition 4.5. *The de Rham cohomology of $\mathrm{Ad}(P)$ is zero except (possibly) in dimension one, where it fits into an exact sequence:*

$$0 \rightarrow f_*(\omega_{X^{\log}/S^{\log}})^{\otimes 2}(-D) \rightarrow \mathbf{R}^1(f_{\mathrm{DR}})_*(\mathrm{Ad}(P)) \rightarrow \mathbf{R}^1 f_* \tau_{X^{\log}/S^{\log}} \rightarrow 0$$

Moreover, (P, ∇_P) has no nontrivial automorphisms.

§4.4. Clutching Morphisms

Let us assume the notation of the subsection of §3 entitled “Clutching Behavior.” Let \mathcal{D} be clutching data for an r -pointed curve of genus g . Let us suppose that for each $i = 1, \dots, n$, we are given a morphism of sets:

$$\mathcal{R}_i : \{1, \dots, r_i\} \rightarrow 0 \cup \Gamma(S, \mathcal{O}_S^\times)$$

with the property that for every edge ϵ with ends ϵ_1 (attached to vertex i) and ϵ_2 (attached to vertex j) such that $\lambda_i(\epsilon_1) = a$ and $\lambda_j(\epsilon_2) = b$, we have $\mathcal{R}_i(a) = \mathcal{R}_j(b)$. Then (in the above notation) whenever we are given a torally indigenous bundle (P_i, ∇_{P_i}) on each $f_i^{\log} : X_i^{\log} \rightarrow S_i^{\log}$ of radii \mathcal{R}_i , it is easy to see that we can glue these torally indigenous bundles together to obtain a torally indigenous bundle (P, ∇_P) on $f^{\log} : X^{\log} \rightarrow S^{\log}$. Indeed, the gluing isomorphisms ξ_ϵ are determined uniquely by the condition that they are compatible with the monodromy operators and Hodge sections on both sides. The radii of (P, ∇_P) are then given by a map of sets

$$\rho : \{1, \dots, r\} \rightarrow 0 \cup \Gamma(S, \mathcal{O}_S^\times)$$

which can be computed combinatorially solely from \mathcal{D} and the $\{\mathcal{R}_i\}$.

Definition 4.6. We shall refer to a collection of such $\{\mathcal{R}_i\}_{i=1, \dots, n}$ as a set of *radii* for \mathcal{D} .

If we carry out this construction in the universal case, we then obtain a commutative diagram of morphisms of moduli stacks:

$$\begin{array}{ccc} \prod_{i=1}^n \overline{\mathcal{S}}_{g_i, r_i}^{\mathcal{R}_i} & \xrightarrow{\kappa_{\mathcal{D}}^\rho} & \overline{\mathcal{S}}_{g, r}^\rho \\ \downarrow & & \downarrow \\ \prod_{i=1}^n \overline{\mathcal{M}}_{g_i, r_i} & \xrightarrow{\kappa_{\mathcal{D}}} & \overline{\mathcal{M}}_{g, r} \end{array}$$

Remark. Let us suppose, for simplicity, that S is the spectrum of an algebraically closed field. Then if we start with a torally indigenous bundle on a nodal curve and restrict to its normalization, we still get a torally indigenous bundle on each component of the normalization. This phenomenon differs from the classical case, for even if one starts with a classical indigenous bundle on a nodal curve, the restriction to the normalization may not have nilpotent monodromy at all the marked points. Thus, by enlarging the category that we work with to include torally indigenous bundles, we can study what happens (for instance) to a classical indigenous bundle on a nodal curve upon restriction to the normalization of the curve. It turns out that this will be of vital importance in the study of degenerations of classical indigenous bundles.

§5. The Universal Torsor of Torally Indigenous Bundles

This § parallels §3 of [Mzk1], Chapter I. Let K be a field of characteristic zero, for instance, \mathbf{Q}_p . Since we are interested only in computing certain intersection numbers, the field is irrelevant. Note that although in the preceding §, we assumed that p is an odd prime, the theory extends immediately to the (much easier) case of characteristic zero. Let $\overline{\mathcal{M}}_{g,r}$ denote the moduli stack of r -pointed stable curves of genus g over K . Let $\rho = \{\rho_i\}$, with $\rho_i \in K$. Then, just as in Theorem 4.4, we get a torsor over $\Omega_{\overline{\mathcal{M}}_{g,r}}^{\log}$

$$\overline{\mathcal{S}}_{g,r}^{\rho} \rightarrow \overline{\mathcal{M}}_{g,r}$$

parametrizing the isomorphism classes of torally indigenous bundles with radii given by ρ on the universal curve. This torsor thus defines a class

$$\Sigma_{g,r}^{\rho} \in H^{1,1}(\overline{\mathcal{M}}_{g,r}^{\log}) \stackrel{\text{def}}{=} H^1(\overline{\mathcal{M}}_{g,r}, \Omega_{\overline{\mathcal{M}}_{g,r}}^{\log})$$

The purpose of this § is to compute this class. It turns out that modulo what was done in [Mzk1], Chapter I, §3 (where $\Sigma_{g,r}^{\rho}$ was computed in the case when all the ρ_i are zero), there is very little work that needs to be done in order to compute this class for arbitrary ρ .

§5.1. Notation

Let us review the notation introduced in [Mzk1], Chapter I, §3. Let $(\zeta^{\log} : \mathcal{C}^{\log} \rightarrow \overline{\mathcal{M}}_{g,r}^{\log}, D)$ be the universal r -pointed stable of genus g ,

i.e., $D \subseteq \mathcal{C}$ is the divisor of marked points. Let $P = \mathbf{P}(\mathcal{I}/\mathcal{I}^{[3]})$ be the \mathbf{P}^1 -bundle of Proposition 4.3. Recall that $P|_D$ is the trivial \mathbf{P}^1 -bundle. Let us denote by \mathcal{C}_c^{\log} the log stack obtained by letting \mathcal{C} be the underlying stack and taking for the log structure the log structure defined by the divisor on \mathcal{C} which is the pull-back (via ζ) of the “divisor at infinity” of $\overline{\mathcal{M}}_{g,r}$. Thus, we have an exact sequence on \mathcal{C} :

$$0 \rightarrow \zeta^* \Omega_{\overline{\mathcal{M}}_{g,r}^{\log}} \rightarrow \Omega_{\mathcal{C}_c^{\log}} \rightarrow \omega_{\mathcal{C}/\overline{\mathcal{M}}_{g,r}} \rightarrow 0$$

where the first two sheaves of differentials are over K . For $i, j \geq 0$, let us define for any $\mathcal{O}_{\mathcal{C}}$ -module \mathcal{F} :

$$H_c^{i,j}(\mathcal{C}, \mathcal{F}) \stackrel{\text{def}}{=} H^i(\mathcal{C}, \mathcal{F} \otimes_{\mathcal{O}_{\mathcal{C}}} \wedge^j \Omega_{\mathcal{C}_c^{\log}})$$

This cohomology is a sort of cohomology with compact supports outside D . Note that the relationship between the logarithmic differentials of \mathcal{C}^{\log} and \mathcal{C}_c^{\log} is given by the following canonical exact sequence:

$$0 \rightarrow \Omega_{\mathcal{C}_c^{\log}} \rightarrow \Omega_{\mathcal{C}^{\log}} \rightarrow \mathcal{O}_D \rightarrow 0$$

where the second arrow is the canonical inclusion, and the third arrow is the residue morphism. By considering the connecting homomorphism of the long exact cohomology sequence associated to this exact sequence of sheaves, we thus obtain a canonical Gysin-type morphism:

$$G : H^0(D, \mathcal{F}|_D) \rightarrow H_c^{1,1}(\mathcal{C}, \mathcal{F})$$

We shall be applying the morphism G to certain canonical sections of $H^0(D, \text{Ad}(P)|_D)$, defined as follows. Let μ_{nilp} be the section of $\text{Ad}(P)|_D$ whose restriction to each σ_i is the “standard monodromy endomorphism of radius zero” (as in Definition 4.2). Let μ_{tor} be the section of $\text{Ad}(P)|_D$ whose restriction to each σ_i is the “standard monodromy endomorphism of radius ρ_i ” (as in Definition 4.2).

Let $\kappa_{\text{zero}} \in H_c^{1,1}(\mathcal{C}, \text{Ad}(P))$ denote the obstruction to putting a connection *relative to the log structure of \mathcal{C}_c^{\log}* (i.e., with zero monodromy at D) on the \mathbf{P}^1 -bundle $P \rightarrow \mathcal{C}$. Let $\kappa_{\text{nilp}} \in H_c^{1,1}(\mathcal{C}, \text{Ad}(P))$ denote the obstruction to putting a connection on $P \rightarrow \mathcal{C}$ relative to the log structure of \mathcal{C}^{\log} with monodromy endomorphism μ_{nilp} at D . Let $\kappa_{\text{tor}} \in H_c^{1,1}(\mathcal{C}, \text{Ad}(P))$ denote the obstruction to putting a connection on $P \rightarrow \mathcal{C}$ relative to the log structure of \mathcal{C}^{\log} with monodromy endomorphism μ_{tor} at D . Note that the κ of [Mzk1], Chapter I, §3, is precisely what we have called κ_{nilp} here. Thus, the key is to relate $\zeta_* \text{tr}(\kappa_{\text{nilp}}^2)$ to $\zeta_* \text{tr}(\kappa_{\text{tor}}^2)$. Symbolically, we can summarize the above definitions as follows:

Lemma 5.1. *We have $\kappa_{\text{nilp}} = \kappa_{\text{zero}} + G(\mu_{\text{nilp}})$; $\kappa_{\text{tor}} = \kappa_{\text{zero}} + G(\mu_{\text{tor}})$; where G is the Gysin-type morphism defined above.*

§5.2. Computation

Now we compute: Note first that if $\mu \in H^0(D, \text{Ad}(P)|_D)$, then we have (by the “projection formula”)

$$\zeta_* \text{tr}(\kappa_{\text{zero}} \cdot G(\mu)) = \gamma_* \text{tr}(\kappa_{\text{zero}}|_D \cdot \mu)$$

where $\gamma : D \rightarrow \overline{\mathcal{M}}_{g,r}$ is the structure morphism. We claim that the right-hand side of this equation is zero: Indeed,

$$\kappa_{\text{zero}}|_D \in H^{1,1}(D, \text{Ad}(P)|_D) \stackrel{\text{def}}{=} H^1(D, \text{Ad}(P)|_D \otimes_{\mathcal{O}_D} \Omega_{\overline{\mathcal{M}}_{g,r}}^{\log}|_D)$$

is the obstruction to putting a connection on the \mathbf{P}^1 -bundle $P|_D$, which is trivial. On the other hand, again, by the projection formula, we have

$$\begin{aligned} \zeta_* \text{tr}(G(\mu)^2) &= \gamma_* \text{tr}(G(\mu)|_D \cdot \mu) \\ &= \gamma_* \text{tr}(\mu^2) \cdot \{([D])|_D\} \\ &= \sum_{i=1}^r (\text{tr}(\mu^2)|_{\sigma_i}) \zeta_* [\sigma_i]^2 \end{aligned}$$

where $[\sigma_i] \in H_c^{1,1}(\mathcal{C}, \mathcal{O}_C)$ is the fundamental class of the divisor σ_i .

Thus, we have

$$\begin{aligned} \zeta_* \text{tr}(\kappa_{\text{nilp}}^2) &= \zeta_* \text{tr}(\kappa_{\text{zero}}^2) + \zeta_* \text{tr}(\kappa_{\text{zero}} \cdot G(\mu_{\text{nilp}})) + \zeta_* \text{tr}(G(\mu_{\text{nilp}})^2) \\ &= \zeta_* \text{tr}(\kappa_{\text{zero}}^2) \end{aligned}$$

and

$$\begin{aligned} \zeta_* \text{tr}(\kappa_{\text{tor}}^2) &= \zeta_* \text{tr}(\kappa_{\text{zero}}^2) + \zeta_* \text{tr}(\kappa_{\text{zero}} \cdot G(\mu_{\text{tor}})) + \zeta_* \text{tr}(G(\mu_{\text{tor}})^2) \\ &= \zeta_* \text{tr}(\kappa_{\text{zero}}^2) + 2 \left(\sum_{i=1}^r \rho_i^2 \zeta_* [\sigma_i]^2 \right) \end{aligned}$$

We summarize this in a Lemma:

Lemma 5.2. *We have: $\zeta_* \text{tr}(\kappa_{\text{tor}}^2) = \zeta_* \text{tr}(\kappa_{\text{nilp}}^2) + 2(\sum_i \rho_i^2 \zeta_* [\sigma_i]^2)$.*

On the other hand, by the same proof as in [Mzk1], Chapter I, Lemma 3.2, we have

Lemma 5.3. *We have: $\Sigma_{g,r}^\rho = \frac{1}{2} \zeta_* \text{tr}(\kappa_{\text{tor}}^2)$.*

Thus, by combining these two Lemmas with [Mzk1], Chapter I, Theorem 3.4, we obtain the following:

Theorem 5.4. *The torsor of isomorphism classes of torally indigenous bundles with radii $\{\rho_i\}$ defines a class $\Sigma_{g,r}^\rho \in H^1(\overline{\mathcal{M}}_{g,r}, \Omega_{\overline{\mathcal{M}}_{g,r}}^{\log})$ which is equal to*

$$\frac{1}{4}(\theta + \psi) - \left(\sum_{i=1}^r \rho_i^2 \psi_i \right)$$

where $\theta = \zeta_* \xi^2$; $\psi_i = \zeta_*(\xi \cdot [\sigma_i])$; $\psi = \sum_i \psi_i$; and $\xi = c_1(\omega_{C/\overline{\mathcal{M}}_{g,r}})$.

Remark. Note, in particular, that because of the extra “negative” term, there is no particular reason why this class should always be nonzero. In fact, we shall soon see a specific example where this class turns out to be zero.

§5.3. The Case of Dimension One

Just as in [Mzk1], Chapter I, §3, in the case of $\dim(\overline{\mathcal{M}}_{g,r}) = 1$, by using cohomology with compact supports, we can get stronger results. Let us concentrate, for simplicity, on the case $g = 1$; $r = 1$. In that case, we can trivialize the torsor $\Sigma_{1,1}^\rho$ at infinity as follows: Let $(E, \sigma : \text{Spec}(K) \rightarrow E)$ be the totally degenerate 1-pointed curve of genus 1. Let $F \rightarrow E$ be the normalization of E . Let $\nu_1, \nu_2 \in F(K)$ be the points which map to the node of E . Let $[E] \in \overline{\mathcal{M}}_{1,1}(K)$ be the point defined by (E, σ) . Then we take as our trivialization of $\Sigma_{1,1}^\rho|_{[E]}$ the unique torally indigenous bundle on (E, σ) of radius ρ whose restriction to F has nilpotent monodromy at ν_1 and ν_2 . This defines a compactly supported class

$$\Sigma_1^\rho \in H^1(\overline{\mathcal{M}}_{1,1}, \Omega_{\overline{\mathcal{M}}_{1,1}/K})$$

Let $\lambda = c_1(\omega_{C/\overline{\mathcal{M}}_{1,1}}|_\sigma) \in H^1(\overline{\mathcal{M}}_{1,1}, \Omega_{\overline{\mathcal{M}}_{1,1}})$. Then

Theorem 5.5. *We have: $\Sigma_1^\rho = (\frac{1}{4} - \rho^2)\lambda$.*

Note, in particular, that this class is zero when $\rho = \pm\frac{1}{2}$. To see why this value of ρ is natural, consider the following. Let (E, σ) be a 1-pointed stable curve of genus 1 over a K -algebra A . Let $(\mathcal{E}, F^1(\mathcal{E}), \nabla_{\mathcal{E}})$ be the canonical classical indigenous bundle on E (i.e., the *zero*-pointed curve E of genus 1) defined in [Mzk1], Chapter I, §2, Example 2. Let $\mathcal{F} \subseteq \mathcal{E}$ be the subbundle of sections whose value at σ lies in the first step $F^1(\mathcal{E})|_{\sigma}$ of the Hodge filtration at σ . Then the connection $\nabla_{\mathcal{E}}$ on \mathcal{E} (which has no pole at σ) induces a connection $\nabla_{\mathcal{F}}$ on \mathcal{F} with a logarithmic pole at σ . Moreover, it is easy to see that $\mathbf{P}(\mathcal{F}, \nabla_{\mathcal{F}})$ defines a torally indigenous bundle on (E, σ) of radius $\pm\frac{1}{2}$. It is this torally indigenous bundle which defines a trivialization of the torsor Σ_1^ρ , for $\rho = \pm\frac{1}{2}$. In summary, we have the following:

Corollary 5.6. *We have: $\overline{\mathcal{S}}_{1,1}^{\pm\frac{1}{2}} \cong \overline{\mathcal{S}}_{1,0}$.*

Chapter II: Torally Crys-Stable Bundles in Positive Characteristic

§0. Introduction

This Chapter is largely of a technical nature: most of the objects and notions that we introduce here will not appear as part of the final theory. The reason we introduce them is that they are technically useful both for proving certain general results (concerning objects that *will* appear as part of the final theory) and for constructing explicit examples of the uniformization theory discussed in this book. In the first § of this Chapter, we study torally crys-stable bundles of positive level in characteristic p whose p -curvature is nilpotent. Moreover, we show (Theorems 1.9, 1.10, and 1.16) that the stack of such bundles satisfies some of the same properties as the stack of nilpotent indigenous bundles (studied in Chapter II, §2, of [Mzk1]). In §2, we study the locus of torally indigenous bundles whose p -curvature is zero in greater detail. In particular, we introduce the notion of a *connection of higher order*, and develop the deformation theory of such connections (Theorem 2.8). Finally, in §3, we study deformations of indigenous bundles whose p -curvature is nilpotent and has zeroes, but is not identically zero (Theorem 3.9).

The common feature of all three sections is that their main purpose is to develop the machinery necessary to solve the following technical “existence” problem: In the rest of this book, we define various kinds of uniformization theories generalizing the ordinary theory of [Mzk1]. Whenever one defines a new kind of uniformization theory, one wants to know that it is nonvacuous, i.e., that there actually *exist* examples of the sort of objects (typically systems of crys-stable bundles satisfying certain conditions) that the uniformization theory in question addresses. Typically, it is too difficult to construct such objects over smooth curves (where the theory most typically takes place). Thus, it is natural to try to show the existence of such data by first constructing such data over, say, a totally degenerate curve

(where, typically, it is relatively easy to construct such data) and then deforming. Indeed, this was precisely the approach taken in [Mzk1] for the “classical” ordinary theory, where we showed (in Chapter II of [Mzk1]) the existence of nilpotent ordinary indigenous bundles by constructing such bundles on totally degenerate curves and then deforming. Unfortunately, the technical problem that typically occurs for the more general uniformization theories discussed in this book is that (unlike in the ordinary case) the sort of data that one wishes to construct (say, ideally, over a smooth curve) simply does not exist over a totally degenerate curve. Thus, we define such notions as *pseudo-torally crys-stable bundles*, *n-connections*, and *mildly spiked bundles*, all of which are rather unnatural and *ad hoc* objects in some sense, but which have the following virtue: Namely, they are relatively easy to construct (and, in particular, *exist!*) over totally degenerate curves, can be deformed to corresponding objects over smooth curves, and, finally, have the property that these corresponding objects over smooth curves are precisely the sort of crys-stable bundles that one wants to construct for the uniformization theory. In Chapter IV, we shall carry out this task of using the theory of the present Chapter to construct examples of data for various uniformization theories.

§1. The p -Curvature of a Torally Crys-Stable Bundle

In this §, we will work entirely in characteristic p (where p is an odd prime). Let S^{\log} be a fine noetherian log scheme of characteristic p . Let $f^{\log} : X^{\log} \rightarrow S^{\log}$ be an r -pointed stable curve of genus g (where $2g - 2 + r \geq 1$). Let $\sigma_1, \dots, \sigma_r : S \rightarrow X$ be the r marked points, and let $D \subseteq X$ be the divisor of marked points. We shall denote by $\Phi_{S^{\log}} : S^{\log} \rightarrow S^{\log}$ the absolute Frobenius on S^{\log} , and by $\Phi_{X^{\log}/S^{\log}} : X^{\log} \rightarrow (X^{\log})^F \stackrel{\text{def}}{=} X^{\log} \times_{S^{\log}, \Phi_{S^{\log}}} S^{\log}$ the relative Frobenius on X^{\log} (over S^{\log}). Suppose that (P, ∇_P) is a *torally crys-stable bundle* on X^{\log} with radii given by

$$\rho : \{1, \dots, r\} \rightarrow 0 \cup \Gamma(S, \mathcal{O}_S^\times)$$

Let $\mathcal{T} \stackrel{\text{def}}{=} \Phi_{X^{\log}/S^{\log}}^* (\tau_{X^{\log}/S^{\log}})^F$. By declaring the sections of $\Phi_{X^{\log}/S^{\log}}^{-1} (\tau_{X^{\log}/S^{\log}})^F$ to be horizontal, we see that we get a natural connection $\nabla_{\mathcal{T}}$ on \mathcal{T} . Recall that the p -curvature

$$\mathcal{P} : \mathcal{T} \rightarrow \text{Ad}(P)$$

of $\text{Ad}(P)$ is a horizontal morphism of vector bundles. The purpose of this § is to study \mathcal{P} , as well as a certain Verschiebung-type morphism (analogous to that of [Mzk1], Chapter II, §2) obtained from \mathcal{P} .

§1.1. Terminology

We begin with some terminology, which will be of use in what follows:

Definition 1.1. We shall call (P, ∇_P) *nilpotent* (respectively, *dormant*; *active*; *admissible*; *torally admissible*) if \mathcal{P} has square nilpotent image (respectively, is identically zero; is not dormant on any fiber of $f : X \rightarrow S$; is injective at every point of X ; is injective away from the nodes and marked points of X^{\log}).

Remark. By “square nilpotent,” we mean the following: Locally sections of $\text{Ad}(P)$ may be regarded as 2-by-2 matrices with trace zero; “square nilpotent” then refers to those matrices whose square is zero.

For instance, we have the following result:

Proposition 1.2. Suppose that $(\pi : P \rightarrow X, \nabla_P)$ is nilpotent and active, and that the pair (g, r) is one of the following: $(0, 3)$; or $(1, 1)$. Then (P, ∇_P) is automatically torally admissible.

Proof. It is easy to see that we may assume that S is the spectrum of an algebraically closed field, and that $f : X \rightarrow S$ is smooth. Suppose that (P, ∇_P) is not torally admissible. Then the p -curvature morphism $\mathcal{P} : \mathcal{T} \rightarrow \text{Ad}(P)$ must vanish to order at least p (since \mathcal{P} is horizontal) at some nonmarked point. Since $\deg(\mathcal{T}) = p$, this implies that \mathcal{P} factors through an injection $\mathcal{M} \rightarrow \text{Ad}(P)$, where \mathcal{M} is a line bundle of degree ≥ 0 . Moreover, the image of this injection is stabilized by the connection on $\text{Ad}(P)$. This contradicts condition (3) of Definition 1.2 of Chapter I. \square

§1.2. The p -Curvature at a Marked Point

In this subsection, we study what happens to the above morphism \mathcal{P} at a marked point σ_i . If we apply $\mathcal{P} : \mathcal{T} = \Phi_{X^{\log}/S^{\log}}^* \tau_{X^{\log}/S^{\log}} \rightarrow \text{Ad}(P)$ to $\Phi_{X^{\log}/S^{\log}}^{-1}(t \frac{d}{dt})$ (where t is a local parameter at the marked point σ_i), and then evaluate the resulting section of $\text{Ad}(\mathcal{E})$ at σ_i , we see that the p -curvature at σ_i is

$$\mu_i^p - \mu_i = \{(\mu_i^2)^{\frac{(p-1)}{2}} - 1\} \cdot \mu_i = (\rho_i^{p-1} - 1) \cdot \mu_i$$

where $\mu_i \in \Gamma(S, \text{Ad}(P)|_{\sigma_i})$ is the monodromy transformation at σ_i , and (since μ_i may be thought of as a two-by-two matrix with trace zero) μ_i^2 is ρ_i^2 times the identity matrix. Now we have two cases: the case where $\rho_i = 0$ (the “classical” case), and the case where ρ_i is a section of \mathcal{O}_S^\times . In a suitable basis, μ_i becomes (the infinitesimal automorphism

of \mathbf{P}^1 induced by) $\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$ (if $\rho_i = 0$), and $\begin{pmatrix} \rho_i & 0 \\ 0 & -\rho_i \end{pmatrix}$ (if $\rho_i \neq 0$). Thus, if $\rho_i = 0$, the p -curvature at σ_i is

$$\begin{pmatrix} 0 & -1 \\ 0 & 0 \end{pmatrix}$$

while for ρ_i invertible, the p -curvature is

$$(\rho_i^p - \rho_i) \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

From this calculation, we can read off a number of results immediately:

Proposition 1.3. *If (P, ∇_P) is nilpotent and admissible, then all the ρ_i are necessarily zero, i.e., (P, ∇_P) is crys-stable.*

Proposition 1.4. *If there exists a ρ_i which is zero, then (P, ∇_P) is active.*

Proposition 1.5. *If (P, ∇_P) is nilpotent, then $(\rho_i^p - \rho_i)^2 = 0$. In particular, if (P, ∇_P) is nilpotent and S is reduced and connected, then $\rho_i \in \mathbf{F}_p$.*

Proposition 1.6. *If all the ρ_i are zero, $f : X \rightarrow S$ is smooth, and (P, ∇_P) is torally admissible, then (P, ∇_P) is, in fact, admissible.*

§1.3. The Verschiebung Morphism

From now on, we assume that all the radii ρ_i are in \mathbf{F}_p . Let us consider the determinant of the p -curvature morphism $\mathcal{P} : \mathcal{T} = \Phi_{X^{\log}/S^{\log}}^* \tau_{X^{\log}/S^{\log}} \rightarrow \text{Ad}(P)$. This determinant naturally defines a horizontal section of $f_*(\mathcal{T}^\vee)^{\otimes 2}$, i.e., a section of $(f_*(\omega_{X^{\log}/S^{\log}})^{\otimes 2})^F$. By the computations in the preceding §, since all the ρ_i are in \mathbf{F}_p , it follows that this section vanishes at the marked points. Thus, we actually obtain a section of $(f_*(\omega_{X^{\log}/S^{\log}})^{\otimes 2}(-D))^F$.

Now suppose that (P, ∇_P) is of a fixed level l . In the universal case, if we define $\overline{\mathcal{Q}}_{g,r}$ to be Spec of the symmetric algebra of $\{(\Omega_{\overline{\mathcal{M}}_{g,r}}^{\log})^\vee\}^F$, i.e., the geometric vector bundle corresponding to $(\Omega_{\overline{\mathcal{M}}_{g,r}}^{\log})^F$, then, just as in [Mzk1], Chapter II, Definition 2.2, we obtain a morphism of algebraic spaces over $\overline{\mathcal{M}}_{g,r}$:

$$\overline{\mathcal{V}}_{g,r}^{\rho_i;l} : \overline{\mathcal{Y}}_{g,r}^{\rho_i;l} \rightarrow \overline{\mathcal{Q}}_{g,r}$$

which we call the *Verschiebung morphism*. We denote by

$$\overline{\mathcal{N}}_{g,r}^{\rho;l} \subseteq \overline{\mathcal{Y}}_{g,r}^{\rho;l}$$

the closed substack which is the inverse image under $\overline{\mathcal{Y}}_{g,r}^{\rho;l}$ of the zero section of $\overline{\mathcal{Q}}_{g,r} \rightarrow \overline{\mathcal{M}}_{g,r}$.

Proposition 1.7. *When $l = \chi$ (i.e., when $(P; \nabla_P)$ is torally indigenous), the morphism $\overline{\mathcal{Y}}_{g,r}^{\rho;l}$ is finite and flat of degree p^{3g-3+r} . In particular, when $l = \chi$, the morphism $\overline{\mathcal{N}}_{g,r}^{\rho;l} \rightarrow \overline{\mathcal{M}}_{g,r}$ is finite and flat of degree p^{3g-3+r} . Moreover, $\overline{\mathcal{N}}_{g,r}^{\rho;l}$ is locally a complete intersection of dimension $3g - 3 + r$.*

Proof. This follows by direct calculation, just as was done in the proof of Theorem 2.3 of [Mzk1], Chapter II. The main point is to show that the Verschiebung morphism $\overline{\mathcal{Y}}_{g,r}^{\rho;l}$, when thought of as a polynomial map between the $(3g-3+r)$ -dimensional affine spaces $\overline{\mathcal{Y}}_{g,r}^{\rho;l} = \overline{\mathcal{S}}_{g,r}^{\rho}$ and $\overline{\mathcal{Q}}_{g,r}$ over $\overline{\mathcal{M}}_{g,r}$, is of degree *precisely* p . Since torally indigenous bundles all have the same underlying \mathbf{P}^1 -bundle (i.e., different torally indigenous bundles on a curve differ only in their connections – cf. Chapter I, §4.1, 4.3), if one starts with a given torally indigenous bundle whose connection is “ ∇ ,” then one may think of the datum of another torally indigenous bundle (on the same curve) as a connection “ $\nabla + \theta$,” where θ is a square nilpotent (i.e., $\theta^2 = 0$) endomorphism-valued differential form. Then one wants to show that *the determinant of the p -curvature of $\nabla + \theta$ is a polynomial in θ of degree precisely p* . But since $\theta^2 = 0$, and the p -curvature may essentially be computed as the p^{th} power of $\nabla + \theta$, it follows that the terms of $(\nabla + \theta)^p$ whose degree (with respect to θ) is maximal are those of the form

$$\theta \cdot \nabla \cdot \dots \nabla \cdot \theta$$

(alternating θ ’s and ∇ ’s, with a total of $\frac{1}{2}(p+1)$ occurrences of θ). That is to say, the degree (with respect to θ) of the p -curvature is $\leq \frac{1}{2}(p+1)$. Now, since one then wants to form the *determinant* of the p -curvature (which essentially amounts to taking the trace of the *square* of the p -curvature), one sees that the degree of the Verschiebung is certainly $\leq 2 \cdot \frac{1}{2}(p+1) = p+1$, and, moreover, the only terms of degree $> p$ will be those that arise as products of two terms of the form just mentioned. But since

$$(\theta \cdot \nabla \cdot \dots \nabla \cdot \theta) \cdot (\theta \cdot \nabla \cdot \dots \nabla \cdot \theta) = 0$$

(again we use the fact that $\theta^2 = 0$), we see that these terms do not contribute. This shows that the Verschiebung is a polynomial of degree $\leq p$ in θ (i.e., in the relative affine variables of $\overline{\mathcal{S}}_{g,r}^{\rho}$ over $\overline{\mathcal{M}}_{g,r}$).

We refer to [Mzk1], Chapter II, Theorem 2.3, for a complete proof of the fact that this degree is exactly p . One heuristic way to convince

oneself of this, however, is to observe (cf. the calculations of §1.2) that since, on a *totally degenerate curve*, the squares $\rho_{\nu,i}^2$ (where $i = 1, \dots, 3g - 3 + r$) of the radii at the nodes (here, the subscript “ ν ” is to distinguish these radii from the radii $\rho = \{\rho_j\}$ at the marked points) form relative affine coordinates for $\overline{S}_{g,r}^\rho$ over $\overline{M}_{g,r}$, the calculations of §1.2 imply that the *Verschiebung in this case* essentially amounts to the morphism

$$\{\rho_{\nu,i}^2\} \mapsto \{(\rho_{\nu,i}^p - \rho_{\nu,i})^2\}$$

(where i ranges from 1 to $3g - 3 + r$) which is clearly polynomial of degree precisely p . \circ

When $l = \chi$, we will denote $\overline{N}_{g,r}^{\rho;l}$ by $\overline{N}_{g,r}^\rho$. Thus, when all the ρ_i are zero, we recover the stack $\overline{N}_{g,r}$ of [Mzk1], Chapter II. Moreover, as g and r vary, these finite, flat morphisms are compatible with all clutching morphisms: that is, for any clutching data \mathcal{D} and radii $\{\mathcal{R}_i\}$ (as in Chapter I, Definitions 3.11 and 4.6), we have natural commutative diagrams:

$$\begin{array}{ccc} \prod_{i=1}^n \overline{N}_{g_i,r_i}^{\mathcal{R}_i} & \longrightarrow & \overline{N}_{g,r}^\rho \\ \downarrow & & \downarrow \\ \prod_{i=1}^n \overline{M}_{g_i,r_i} & \xrightarrow{\kappa_{\mathcal{D}}} & \overline{M}_{g,r} \end{array}$$

Let $k = \mathbf{F}_p$; $S = \mathrm{Spec}(k)$; endow S with the log structure associated to the morphism of monoids $\mathbf{N} \rightarrow \mathcal{O}_S$ given by $1 \mapsto 0$; call the resulting log scheme S^{\log} . Let $X^{\log} \rightarrow S^{\log}$ be a *totally degenerate r -pointed stable curve of genus g* . Let $\delta \in \overline{M}_{g,r}(k)$ be the point corresponding to $X^{\log} \rightarrow S^{\log}$. Let $\delta^\rho \in \overline{N}_{g,r}^\rho(k)$ be a point lying over δ . Thus, δ^ρ corresponds to some torally indigenous bundle ζ on X^{\log} obtained by gluing together various torally indigenous bundles ζ_C on the irreducible components C of X . Let s_C be the number of nonzero radii of ζ_C ; and let

$$s = \frac{1}{2} \sum_C s_C$$

Then the explicit computation of the p -curvature (and hence, by taking its determinant, of the *Verschiebung*) in the preceding subsection shows that the fiber of $\overline{N}_{g,r}^\rho$ over $\overline{M}_{g,r}$ at δ^ρ is (noncanonically) isomorphic to

$$\mathrm{Spec}\left(\bigotimes_{i=1}^s k[\epsilon_i]/(\epsilon_i^2)\right)$$

In particular, when $s = 0$, $\overline{N}_{g,r}^\rho$ is *étale* over $\overline{M}_{g,r}$ at δ^ρ . This result is consistent with the result of [Mzk1] (i.e., [Mzk1], Chapter II, §3,

Proposition 3.7) to the effect that if all the $\rho_i = 0$ ($i = 1, \dots, r$), and $s = 0$ (so that ζ is an indigenous bundle of *restricted type* in the terminology of [Mzk1]), then $\overline{\mathcal{N}}_{g,r}^\rho = \overline{\mathcal{N}}_{g,r}$ is étale over $\overline{\mathcal{M}}_{g,r}$ at δ^ρ . At any rate, we have the following result:

Proposition 1.8. *In the above notation, the complete local ring $\widehat{\mathcal{O}}_{\overline{\mathcal{N}}_{g,r}, \delta^\rho}^\rho$ is finite and flat over $\widehat{\mathcal{O}}_{\overline{\mathcal{M}}_{g,r}, \delta}$ of degree $2^s \leq 2^{3g-3+r}$. In particular, if $p > 2^{3g-3+r}$, then every irreducible component of $(\overline{\mathcal{N}}_{g,r}^\rho)_{\text{red}}$ is generically étale over $\overline{\mathcal{M}}_{g,r}$.*

Remark. The latter half of the above Proposition will be strengthened in Chapter V (cf. Theorem 1.1 and Corollary 1.2 of Chapter V).

§1.4. Torally Crys-Stable Bundles of Arbitrary Positive Level

We would like to generalize (as much as is possible of) Proposition 1.7 to the case of arbitrary positive l . First note that by restricting the natural morphism $\Delta_{g,r}^l : \overline{\mathcal{Y}}_{g,r}^{\rho;l} \rightarrow \overline{\mathcal{D}}_{g,r}^l$ (cf. the discussion following Chapter I, Definition 3.6), we obtain a morphism $\overline{\mathcal{N}}_{g,r}^{\rho;l} \rightarrow \overline{\mathcal{D}}_{g,r}^l$. Observe that $\overline{\mathcal{N}}_{g,r}^{\rho;l} \subseteq \overline{\mathcal{Y}}_{g,r}^{\rho;l}$ is a closed substack defined locally by $3g - 3 + r$ equations, and that (by Lemma 3.8 of Chapter I) we have $\dim_{\mathbf{F}_p}(\overline{\mathcal{D}}_{g,r}^l) + (3g - 3 + r) = \dim_{\mathbf{F}_p}(\overline{\mathcal{Y}}_{g,r}^{\rho;l})$. Let $(\overline{\mathcal{D}}_{g,r}^l)' \subseteq \overline{\mathcal{D}}_{g,r}^l$ denote the open substack whose k -valued points (for k an algebraically closed field) are those l -divisors balanced divisors $D \subseteq X$ (on a stable, pointed curve X) that satisfy the following property:

(*)The multiplicity in the divisor D of every point $x \in X(k)$ is $\leq p - 2$.

Let $(\overline{\mathcal{N}}_{g,r}^{\rho;l})' \subseteq \overline{\mathcal{N}}_{g,r}^{\rho;l}$ be the open substack which is the inverse image of $(\overline{\mathcal{D}}_{g,r}^l)'$ under the morphism $\overline{\mathcal{N}}_{g,r}^{\rho;l} \rightarrow \overline{\mathcal{D}}_{g,r}^l$.

Now let us consider the following useful construction:

Indigenization of a Torally Crys-Stable Bundle of Positive Level: Let k be a field. Let X^{\log} be an r -pointed stable log-curve of genus g over $\text{Spec}(k)$ (equipped with some log structure). Suppose that $(P; \nabla_P)$ is a nilpotent torally crys-stable bundle of level l on X^{\log} , with Kodaira-Spencer locus $D_\kappa \subseteq X$. Thus, D_κ avoids the marked points and nodes of X , and (after possibly enlarging k) can be written in the form $\sum_{i=1}^a e_i \cdot p_i$, where the p_i are k -valued smooth non-marked points of X . Let us write $P = \mathbf{P}(\mathcal{E})$, where \mathcal{E} is a vector bundle on X of rank two. Let $\mathcal{L} \subseteq \mathcal{E}$ be the subbundle of rank one corresponding to the section of $P \rightarrow X$ of negative canonical height. We would like to define a new vector bundle \mathcal{F} such that $\mathcal{E} \subseteq \mathcal{F}$ as follows: Let t_i be a local coordinate

which is equal to 0 at p_i . Then we take \mathcal{F} to be the inductive limit of the following diagram:

$$\begin{array}{ccc} \mathcal{L} & \longrightarrow & \mathcal{E} \\ \downarrow t_i^{e_i+1} & & \\ \mathcal{L} & & \end{array}$$

in a neighborhood of p_i , and we take \mathcal{F} to be equal to \mathcal{E} away from the p_i . Let $Q = \mathbf{P}(\mathcal{F})$. Let Y^{\log} be the pointed stable log-curve obtained from X^{\log} by adding (as marked points of Y^{\log}) the p_i . Then one checks easily that ∇_P defines a logarithmic connection ∇_Q on $Q \rightarrow Y$ (with respect to the log structure of Y^{\log}), and that the locally split injection $\mathcal{L} \hookrightarrow \mathcal{E}$ defines a locally split injection $\mathcal{M} \subseteq \mathcal{F}$ (where \mathcal{M} is a line bundle). Moreover, the Kodaira-Spencer morphism of the filtration $\mathcal{M} \subseteq \mathcal{F}$ is easily seen to be an isomorphism (by the construction of \mathcal{F}). Thus, $(Q; \nabla_Q)$ is a *torally indigenous bundle* on Y^{\log} of radius $\frac{1}{2}(e_i+1) \pmod{p}$ at p_i . We shall refer to (Q, ∇_Q) as the *indigenization* of (P, ∇_P) .

Note that in order to reconstruct (P, ∇_P) from (Q, ∇_Q) , it suffices to specify the $\{\mathcal{O}_X/(t_i^{e_i+1})\}$ -flat submodules $\mathcal{K}_i \subseteq \mathcal{F}/(t_i^{e_i+1} \cdot \mathcal{F})$ generated by the image of \mathcal{E} . Note that this submodule is stabilized by the projective connection ∇_Q . Thus, if $e_i+1 \leq p-1$ (which is the case if the divisor D_κ lies in $(\overline{\mathcal{D}}_{g,r}^l)'$), \mathcal{K}_i is completely determined by the subspace $\mathcal{K}_i \otimes \mathcal{O}_X/(t_i)$ of $\mathcal{F} \otimes \mathcal{O}_X/(t_i)$. But this subspace is completely determined as the unique subspace of $\mathcal{F} \otimes \mathcal{O}_X/(t_i)$ fixed by the monodromy operator of (Q, ∇_Q) at p_i with some eigenvalue (determined by e_i). Thus, in summary, if the divisor D_κ lies in $(\overline{\mathcal{D}}_{g,r}^l)'$, then the correspondence $(P, \nabla_P) \mapsto (Q, \nabla_Q)$ given by taking the indigenization is injective. This completes our discussion of the indigenization of a torally crys-stable bundle of positive level.

Theorem 1.9. *Suppose that $l > 0$. Then the morphism $(\overline{\mathcal{N}}_{g,r}^{\rho;l})' \rightarrow (\overline{\mathcal{D}}_{g,r}^l)'$ is quasi-finite and flat. Moreover, $(\overline{\mathcal{N}}_{g,r}^{\rho;l})'$ is locally a complete intersection of dimension $3g-3+r+2\chi-2l$ over \mathbf{F}_p .*

Proof. By the dimension count reviewed above, it suffices to prove that the morphism $(\overline{\mathcal{N}}_{g,r}^{\rho;l})' \rightarrow (\overline{\mathcal{D}}_{g,r}^l)'$ is quasi-finite. Thus, we fix an algebraically closed field k , an r -pointed stable log-curve X^{\log} of genus g over k , and an l -balanced divisor $D \subseteq X$ that lies in $(\overline{\mathcal{D}}_{g,r}^l)'$. We wish to show that the number of torally crys-stable bundles on X^{\log} of level l with Kodaira-Spencer locus D is finite. But since the divisor D lies in $(\overline{\mathcal{D}}_{g,r}^l)'$, passing to the indigenization is an injective operation. Moreover, by Proposition 1.7, the number of possible indigenizations is finite. This completes the proof. \circ

Although, unfortunately, we cannot show that the morphism $(\overline{\mathcal{N}}_{g,r}^{\rho;l})' \rightarrow (\overline{\mathcal{D}}_{g,r}^l)'$ is *finite*, we can nevertheless show that the following

form of properness is satisfied. Let $S = \operatorname{Spec}(A)$, where A is a discrete valuation ring of characteristic p . Let S^{\log} be a log scheme whose underlying scheme is S and whose log structure is defined by the divisor constituted by the closed point. Let η be the generic point of S . Recall that we denote by $\overline{\mathcal{Y}}_{g,r}^{\rho}$ the stack of torally crys-stable bundles of radii ρ .

Theorem 1.10. *Let $l \geq 0$. Suppose that we are given a morphism $\beta_{\eta} : \eta \rightarrow \overline{\mathcal{N}}_{g,r}^{\rho;l}$ whose composite with $\overline{\mathcal{N}}_{g,r}^{\rho;l} \rightarrow \overline{\mathcal{D}}_{g,r}^l$ extends to a morphism $S \rightarrow \overline{\mathcal{D}}_{g,r}^l$. Then the composite of β_{η} with the inclusion $\overline{\mathcal{N}}_{g,r}^{\rho;l} \subseteq \overline{\mathcal{Y}}_{g,r}^{\rho}$ extends to a morphism $S \rightarrow \overline{\mathcal{Y}}_{g,r}^{\rho}$.*

Proof. We suppose that we are given data as follows: Let $X^{\log} \rightarrow S^{\log}$ be an r -pointed stable log-curve of genus g which is smooth over the generic point η of S . Let $(P_{\eta}, \nabla_{P_{\eta}})$ be a *nilpotent* torally crys-stable bundle of level $l > 0$ on X_{η}^{\log} . By enlarging A , we may assume that the divisor where the Kodaira-Spencer morphism D_{η} of this crys-stable bundle vanishes is a union of sections of $X_{\eta}^{\log} \rightarrow \eta$. Let $(Q_{\eta}, \nabla_{Q_{\eta}})$ be the *indigenization* of $(P_{\eta}, \nabla_{P_{\eta}})$. Thus, $(Q_{\eta}, \nabla_{Q_{\eta}})$ is a torally indigenous bundle on the curve Y_{η}^{\log} , where $Y^{\log} \rightarrow S^{\log}$ is an $(r+r')$ -pointed (where $r' \leq (2\chi - 2l)$ is the number of marked points of Y^{\log} that are not marked points of X^{\log}) stable log-curve of genus g obtained by adding the sections of $X_{\eta}^{\log} \rightarrow \eta$ that make up D_{η} as marking sections to the marking sections of X_{η}^{\log} . Note that $(Q_{\eta}, \nabla_{Q_{\eta}})$ is still nilpotent, hence by Proposition 1.7, admits an extension (Q, ∇_Q) to Y^{\log} .

Let $D \subseteq X$ be the S -flat closure of D_{η} in X . By assumption, D is l -balanced; in particular, it avoids the nodes and marked points of X^{\log} . Let $D_s \subseteq D$ be the restriction to the special fiber of D . Let $U \subseteq X$ be the complement of D_s in X . Then U also defines an open subscheme of Y . Thus, by restricting (Q, ∇_Q) to U (regarded as an open subscheme of Y), we get a \mathbf{P}^1 -bundle with connection $(Q, \nabla_Q)|_U$. Let $P_U \rightarrow U$ be the \mathbf{P}^1 -bundle defined as follows: over U_{η} , we let $P_U \rightarrow U$ be the restriction of $P_{\eta} \rightarrow X_{\eta}$ to U_{η} ; over $X - D$, we let $P_U \rightarrow U$ be the restriction of $(Q, \nabla_Q)|_U$ to $X - D$. Since $U = (X - D) \cup X_{\eta}$, this completes the definition of $P_U \rightarrow U$. Also, it is clear that the connections $\nabla_{P_{\eta}}$ and $\nabla_{Q|_U}$ define a connection ∇_{P_U} on P_U with respect to the restriction of the log structure of X^{\log} to U . On the other hand, since X is *regular of dimension 2* at the points of D_s , it follows from basic commutative algebra (namely, the Auslander-Buchsbaum formula) that the \mathbf{P}^1 -bundle $P_U \rightarrow U$ extends to a \mathbf{P}^1 -bundle $P \rightarrow X$ over X . (The reader who dislikes extending \mathbf{P}^1 -bundles can think in terms of extending the *vector bundle* $\operatorname{Ad}(P_U)$.) Moreover, the connection ∇_{P_U} extends immediately to a connection ∇_P on $P \rightarrow X$. Hence, we obtain a \mathbf{P}^1 -bundle with connection (P, ∇_P) on X^{\log} . Moreover, we are now in the situation of Proposition 3.9 of Chapter I, and so we conclude that (P, ∇_P) is torally crys-stable on X^{\log} . This completes the the proof of the Theorem. \circ

§1.5. The Geometric Connectedness of $\overline{\mathcal{N}}_{g,r}^\rho$

We begin with a Lemma:

Lemma 1.11. *Suppose that the highest Chern class $c_{3g-3+r}(\Omega_{\overline{\mathcal{M}}_{g,r}}^{\log})$ is nonzero. Then any two irreducible components of $\overline{\mathcal{N}}_{g,r}^\rho \otimes_{\mathbf{F}_p} \overline{\mathbf{F}}_p$ intersect. In particular, $\overline{\mathcal{N}}_{g,r}^\rho$ is geometrically connected.*

Proof. Let us work over $\overline{\mathbf{F}}_p$, the algebraic closure of \mathbf{F}_p . Let I_1 and I_2 be any two non-intersecting irreducible components I_1 and I_2 of $\overline{\mathcal{N}}_{g,r}^\rho$. Let us regard I_1 and I_2 as objects over $\overline{\mathcal{M}}_{g,r}$. Then one can always choose a projective variety V that dominates I_1 and I_2 as an object over $\overline{\mathcal{M}}_{g,r}$, i.e., fits into a commutative diagram of generically finite morphisms as follows:

$$\begin{array}{ccc} V & \longrightarrow & I_1 \\ \downarrow & & \downarrow \\ I_2 & \longrightarrow & \overline{\mathcal{M}}_{g,r} \end{array}$$

Moreover, since $I_j \subseteq \overline{\mathcal{N}}_{g,r}^\rho \subseteq \overline{\mathcal{S}}_{g,r}^\rho$ (for $j = 1, 2$), we obtain morphisms $\psi_j : V \rightarrow \overline{\mathcal{S}}_{g,r}^\rho$ (for $j = 1, 2$) over $\overline{\mathcal{M}}_{g,r}$ such that $\psi_1(x) \neq \psi_2(x)$, for every $x \in V$. Let \mathcal{G} be the pull-back of the vector bundle $\Omega_{\overline{\mathcal{M}}_{g,r}}^{\log}$ to V . Then since $\overline{\mathcal{S}}_{g,r}^\rho$ is a torsor over $\Omega_{\overline{\mathcal{M}}_{g,r}}^{\log}$, by subtracting ψ_1 from ψ_2 , we obtain a section $s \in \Gamma(V, \mathcal{G})$ which is *nonzero everywhere*. But by basic intersection theory (using the Chow ring, as in, say, [Fulton1]), this contradicts the fact that $c_{3g-3+r}(\mathcal{G}) \neq 0$. \circ

Remark. In fact, it follows from the results of [HZ] that the hypotheses of the above Lemma are satisfied. Indeed, it follows from the “logarithmic Chern-Gauss-Bonnet Theorem” that the Chern class in the Lemma is equal to the “orbifold Euler characteristic of $\mathcal{M}_{g,r}$ ” (in the terminology of [HZ]). (Note that the logarithmic Chern-Gauss-Bonnet Theorem can be derived immediately from the usual Chern-Gauss-Bonnet Theorem by applying the usual Chern-Gauss-Bonnet Theorem to coverings of $\overline{\mathcal{M}}_{g,r}$ that are ramified to high order at the divisor at infinity, and taking the limit as the ramification index goes to infinity.) Moreover, the results of [HZ] imply that the “orbifold Euler characteristic of $\mathcal{M}_{g,r}$ ” is nonzero. Thus, *we conclude, in particular, that any two irreducible components of $\overline{\mathcal{N}}_{g,r}^\rho \otimes_{\mathbf{F}_p} \overline{\mathbf{F}}_p$ intersect* (see Theorem 1.12 below). Nevertheless, since [HZ] involves deep techniques that lie beyond the scope of this book, it is interesting to see to what extent one can prove that $\overline{\mathcal{N}}_{g,r}^\rho$ is geometrically connected within the context of the techniques of this book. We proceed to do this in the following paragraphs. Another reason why we give an elementary proof of the

weaker connectedness result is that perhaps refinement of the technique discussed below can lead to an alternate proof (in the context of the present book) of the fact that the Chern class of Lemma 1.11 is nonzero.

At any rate, at least when $\dim(\overline{\mathcal{M}}_{g,r}) = 1$, i.e., when the pair (g, r) is $(0, 4)$ or $(1, 1)$, there is no problem in checking that the hypotheses of Lemma 1.11 are satisfied. Now we apply induction on the dimension of $\overline{\mathcal{M}}_{g,r}$ to the statement:

$(*_g, r) \overline{\mathcal{N}}_{g,r}^\rho$ is geometrically connected for all ρ .

We know that $(*_g, r)$ is true when $\overline{\mathcal{M}}_{g,r}$ has dimension one, by Lemma 1.11.

The induction step is proven as follows. To fix ideas, let us assume that $g \geq 1$. (The case when $g = 0$ is proven similarly, and differs only combinatorially from the case considered.) Note that we have a clutching morphism

$$\kappa : \overline{\mathcal{M}}_{1,1} \times \overline{\mathcal{M}}_{g-1,r+1} \rightarrow \overline{\mathcal{M}}_{g,r}$$

obtained by gluing the genus 1 curve at its sole marked point to the genus $g - 1$ curve at its $(r + 1)^{\text{th}}$ marked point. The image of κ is a divisor at infinity of $\overline{\mathcal{M}}_{g,r}$. Let us consider what $\overline{\mathcal{N}}_{g,r}^\rho$ looks like over the image of κ . This amounts to considering nilpotent torally indigenous bundles (P, ∇_P) on an r -pointed curve $f^{\log} : X^{\log} \rightarrow S^{\log}$ of genus g which is obtained by gluing together a 1-pointed curve $(X')^{\log} \rightarrow S^{\log}$ of genus 1 to an $(r + 1)$ -pointed curve $(X'')^{\log} \rightarrow S^{\log}$ of genus $g - 1$. Since we are only interested in connectedness, we may assume that S is reduced. Let us also assume that S is *geometrically connected*. If we restrict (P, ∇_P) to $(X')^{\log}$ (respectively, $(X'')^{\log}$), we obtain a nilpotent torally indigenous bundle $(P', \nabla_{P'})$ (respectively, $(P'', \nabla_{P''})$) with radii given by some set ρ' (respectively, ρ'') of elements of \mathbf{F}_p (by Proposition 1.5). In other words, over κ , we have a clutching morphism of nilpotent torally indigenous bundles

$$\kappa^{\rho', \rho''} : \overline{\mathcal{N}}_{1,1}^{\rho'} \times \overline{\mathcal{N}}_{g-1,r+1}^{\rho''} \rightarrow \overline{\mathcal{N}}_{g,r}^\rho$$

By the induction hypothesis, the image of $\kappa^{\rho', \rho''}$ is geometrically connected. Moreover, as ρ' and ρ'' range over all possibilities, the union of the images of the various $\kappa^{\rho', \rho''}$ is all of $\overline{\mathcal{N}}_{g,r}^\rho \times_{\overline{\mathcal{M}}_{g,r}; \kappa} (\overline{\mathcal{M}}_{1,1} \times \overline{\mathcal{M}}_{g-1,r+1})$. Thus, to complete the proof of the induction step, it suffices to prove that the images of different $\kappa^{\rho', \rho''}$ belong to the same connected component of $\overline{\mathcal{N}}_{g,r}^\rho$.

Note that the radii that make up ρ' or ρ'' are completely determined by ρ , *except for the one radius ρ_c* (which is the same in the primed and double primed cases) at the marked point where the clutching

takes place. (We shall refer to ρ_c as the *clutching radius*.) Thus, for simplicity, we shall often write κ^{ρ_c} for $\kappa^{\rho', \rho''}$. Next, we consider the (rather trivial) clutching morphism $\overline{\mathcal{M}}_{0,3} \rightarrow \overline{\mathcal{M}}_{1,1}$. Taking the product of this morphism with $\overline{\mathcal{M}}_{g-1,r+1}$ and then composing this with the previous clutching morphism, we obtain a morphism

$$\lambda : \overline{\mathcal{M}}_{0,3} \times \overline{\mathcal{M}}_{g-1,r+1} \rightarrow \overline{\mathcal{M}}_{g,r}$$

Observe that this clutching morphism λ also factors through the clutching morphism $\xi : \overline{\mathcal{M}}_{g-1,r+2} \rightarrow \overline{\mathcal{M}}_{g,r}$. Thus, in summary, we have a commutative diagram:

$$\begin{array}{ccc} \overline{\mathcal{M}}_{0,3} \times \overline{\mathcal{M}}_{g-1,r+1} & \longrightarrow & \overline{\mathcal{M}}_{1,1} \times \overline{\mathcal{M}}_{g-1,r+1} \\ \downarrow & & \downarrow \kappa \\ \overline{\mathcal{M}}_{g-1,r+2} & \xrightarrow{\xi} & \overline{\mathcal{M}}_{g,r} \end{array}$$

where both composites are equal to λ .

Now, let us start with a curve Y^{\log} in the image of λ and a nilpotent torally indigenous bundle (Q, ∇_Q) on it (with radii given by ρ), whose clutching radius on the left (i.e., the $(0, 3)$ -curve) is ρ_l and on the right (i.e., the $(g-1, r+1)$ -curve) is ρ_r . Then let us take another curve Z^{\log} in the image of λ . Let us also take a nilpotent torally indigenous bundle (W, ∇_W) on it (with radii given by ρ), whose clutching radius on the left is some ρ_l and whose clutching radius on the right is $\rho_{r'}$ (possibly distinct from ρ_r). Note that both the pair $(Y^{\log}; (Q, \nabla_Q))$ and the pair $(Z^{\log}; (W, \nabla_W))$ define points of $\overline{\mathcal{N}}_{g,r}^{\rho}$ that are in the image of the clutching morphism

$$\xi^{\rho_l} : \overline{\mathcal{N}}_{g-1,r+2}^{\rho'''} \rightarrow \overline{\mathcal{N}}_{g,r}^{\rho}$$

for some appropriate ρ''' (which depends on ρ and ρ_l , but *not* on ρ_r or $\rho_{r'}$). But by the induction hypothesis, the image of ξ^{ρ_l} is connected. This completes the proof of the induction step. In summary, we have given a complete elementary proof of the second statement of the following Theorem:

Theorem 1.12. *Any two irreducible components of $\overline{\mathcal{N}}_{g,r}^{\rho} \otimes_{\mathbf{F}_p} \overline{\mathbf{F}}_p$ intersect. In particular, $\overline{\mathcal{N}}_{g,r}^{\rho}$ is geometrically connected.*

Remark. In the case $g = 0$, we use the following diagram of clutching morphisms:

$$\begin{array}{ccc} \overline{\mathcal{M}}_{0,r-2} \times \overline{\mathcal{M}}_{0,3} \times \overline{\mathcal{M}}_{0,3} & \longrightarrow & \overline{\mathcal{M}}_{0,r-1} \times \overline{\mathcal{M}}_{0,3} \\ \downarrow & & \downarrow \kappa \\ \overline{\mathcal{M}}_{0,r-2} \times \overline{\mathcal{M}}_{0,4} & \xrightarrow{\xi} & \overline{\mathcal{M}}_{g,r} \end{array}$$

This completes the induction for the case ignored above.

Remark. Note that even if one is only interested in the result for the classical case (when all the ρ_i are zero), in order to perform the induction, one needs to make essential use of the $\overline{\mathcal{N}}_{g,r}^\rho$ for all possible ρ .

Remark. It is tempting to conjecture further that $\overline{\mathcal{N}}_{g,r}^\rho$ is, in fact, irreducible, at least perhaps when $\rho = 0$ (the classical case). As we shall see in Chapter IV, §1, however, this is not necessarily the case.

§1.6. Degenerations of Torally Crys-Stable Bundles of Positive Level

In this subsection, we define stacks of nilpotent “PTCS” (or “pseudo-torally crys-stable”) bundles. The definition of a PTCS bundle is essentially the same as that of a crys-stable bundle of positive level when the Kodaira-Spencer locus is *étale* over the base, but differs more and more, as the Kodaira-Spencer locus becomes ramified over the base. (In fact, for the sake of simplicity, we shall restrict ourselves here to the simplest case, that is, the case where the Kodaira-Spencer locus is a divisor of degree 2.) Although these PTCS bundles are not very useful from the point of view of uniformization theory, they are useful for studying how torally crys-stable bundles degenerate as the underlying curve degenerates. For instance, on a totally degenerate curve, it is not difficult to see that the only torally crys-stable bundles of positive level are torally indigenous. Thus, by restricting ourselves to the “torally crys-stable category,” it is difficult to see what happens as non-indigenous positive level torally crys-stable bundles degenerate. Also, just as in the case of dormant torally crys-stable bundles (studied in §2 of this Chapter), since it is difficult to construct nilpotent positive level torally crys-stable bundles directly on smooth curves, it is convenient to be able to construct some sort of natural generalization of such bundles on stable curves. Then by deforming these bundles to bundles on smooth curves, we can conclude the existence of such bundles on smooth curves.

Let S^{\log} be a fine noetherian log scheme of characteristic $p \geq 5$. Let $Y^{\log} \rightarrow S^{\log}$ be an $(r+2)$ -pointed stable curve of genus g (where $2g-2+r \geq 1$). Let ρ be a set of r radii in \mathbf{F}_p . We shall refer to the last 2 marked points of Y^{\log} as *auxiliary*. We shall think of the radii of ρ as corresponding to the first r , or *nonauxiliary*, marked points of Y^{\log} . Let ρ' be the set of $r+2$ radii whose first r radii are those given by ρ , and whose last 2 radii are equal to 1. Let (Q, ∇_Q) be a \mathbf{P}^1 -bundle with connection on Y^{\log} of radii ρ' . We would like to define the *adjustment of (Q, ∇_Q) at an auxiliary marked point* as follows: Let $\sigma : S \rightarrow Y$ be an auxiliary marked point. Since the construction will be *étale* local on S and canonical, we are always free to replace S by some S' which is *étale* and surjective over S . Let $U \subseteq Y$ be a neighborhood of the image of σ such that the restriction of (Q, ∇_Q) to U is the projectivization of a

rank two bundle with connection $(\mathcal{E}, \nabla_{\mathcal{E}})$ whose determinant is trivial. Let $\mathcal{I} \subseteq \mathcal{O}_U$ be the ideal defining the image of σ . Let $Z = V(\mathcal{I}^2)$. Thus, $Z \subseteq U$ is the first infinitesimal neighborhood of σ . Let $\mathcal{K} \subseteq \mathcal{E}|_Z$ be the unique invertible \mathcal{O}_Z -submodule of $\mathcal{E}|_Z$ such that the monodromy action on $\sigma^*\mathcal{K}$ has eigenvalue 1, and \mathcal{K} is fixed by $\nabla_{\mathcal{E}}$. Let $\mathcal{F}' \subseteq \mathcal{E}$ be the sub- \mathcal{O}_U -module of sections whose restriction to Z lies in \mathcal{K} . Let $\mathcal{F} = \mathcal{F}' \otimes_{\mathcal{O}_U} \mathcal{I}^{-1}$. Note that $\nabla_{\mathcal{E}}$ induces a natural logarithmic connection $\nabla_{\mathcal{F}}$ on \mathcal{F} . Let (P, ∇_P) be the \mathbf{P}^1 -bundle with connection on Y^{\log} which is equal to (Q, ∇_Q) away from σ and equal to the projectivization of $(\mathcal{F}, \nabla_{\mathcal{F}})$ over U . Note that $(\mathcal{F}, \nabla_{\mathcal{F}})$ has square nilpotent monodromy at σ . Let \mathcal{L} be the line bundle $\mathcal{K}^{-1}|_{\text{Im}(\sigma)} \otimes_{\mathcal{O}_S} \sigma^*\mathcal{I}$ on S . Observe that the image of $\mathcal{E} \otimes_{\mathcal{O}_U} \mathcal{I}$ in \mathcal{F} defines a subbundle of $\sigma^*\mathcal{F}$ which is isomorphic to \mathcal{L} and which is fixed by the monodromy operator arising from $\nabla_{\mathcal{F}}$ at σ . That is to say, we have a filtration $\mathcal{L} \subseteq \sigma^*\mathcal{F}$ (such that $\sigma^*\mathcal{F}/\mathcal{L} \cong \mathcal{L}^{\vee}$) with respect to which the monodromy operator of $\nabla_{\mathcal{F}}$ acts in a nilpotent fashion. In particular, the monodromy of $(\mathcal{F}, \nabla_{\mathcal{F}})$ at σ defines a section μ_{σ}^{ad} of $\mathcal{L}^{\otimes 2}$ over S . Moreover, the monodromy of $(\mathcal{F}, \nabla_{\mathcal{F}})$ at σ is zero if and only if μ_{σ}^{ad} is identically zero.

Definition 1.13. We shall refer to (P, ∇_P) as the *adjustment* of (Q, ∇_Q) at the *auxiliary marked section* σ . We shall refer to the closed subscheme of S given by $V(\mu_{\sigma}^{\text{ad}})$ (i.e., the zero locus of μ_{σ}^{ad}) as the *adjustable locus* of (Q, ∇_Q) for σ . We shall say that (Q, ∇_Q) is *adjustable* at σ if μ_{σ}^{ad} is identically zero on S .

Let $S = \overline{\mathcal{S}}_{g, r+2}^{\rho'}$. Let $Y^{\log} \rightarrow S^{\log}$ be the pull-back of the tautological log-curve over $\overline{\mathcal{M}}_{g, r+2}^{\log}$. Let (Q, ∇_Q) be the tautological torally indige-nous bundle of radii ρ' on Y^{\log} . Let $\mathcal{N}' \subseteq \overline{\mathcal{S}}_{g, r+2}^{\rho'}$ be the adjustable locus for the last marking section. Endow \mathcal{N}' with the pull-back to \mathcal{N}' of the log structure of $\overline{\mathcal{M}}_{g, r+2}^{\log}$. Let $(Y' \rightarrow \mathcal{N}', M_{Y'})$ be the $(r+1)$ -pointed semi-stable curve of genus g (where $M_{Y'} \subseteq Y'$ is the divisor of marked points) obtained by forgetting the last marking section of Y^{\log} and then restricting to \mathcal{N}' . Let $(Y')^{\log}$ be the log structure on Y' which co-incides with that of $Y^{\log}|_{\mathcal{N}'}$ away from the last marking section, and is the pull-back of the log structure of \mathcal{N}' near the last marking section. Let $Y' \rightarrow Z$ be the smallest contraction such that Z together with the images of the nonauxiliary marking sections of Y' is *semi-stable* as an r -pointed curve of genus g . Note that the image of the auxiliary section of Y' avoids the nodes of Z , as well as the image in Z of $M_{Y'}$. (In other words, in a fiber Y'_ν of Y' over \mathcal{N}' , the only sorts of \mathbf{P}^1 's that we contract are those with two auxiliary marked points and only one other node/marked point.) Thus, we obtain a semi-stable $(r+1)$ -pointed curve $(Z \rightarrow \mathcal{N}', M_Z)$ of genus g (where M_Z is the image in Z of $M_{Y'}$). Let $\sigma_Z : \mathcal{N}' \rightarrow Z$ be the composite of the auxiliary marking section $\mathcal{N}' \rightarrow Y'$ of Y' with the contraction $Y' \rightarrow Z$. Then over $W \stackrel{\text{def}}{=} Z - \text{Im}(\sigma_Z)$, $Y' \rightarrow Z$ is an isomorphism. Let us endow Z with the log structure obtained by taking the log structure of Y' over W , and taking the natural “marked point log structure” at the marking section σ_Z (see, e.g., [Mzk1], Chapter I, Definition 1.1). Then $Y' \rightarrow Z$ extends to a

logarithmic morphism $(Y')^{\log} \rightarrow Z^{\log}$. Let $V \subseteq Z$ be the largest open set over which $Y' \rightarrow Z$ is an isomorphism. Note that $W \subseteq V$, and, moreover, by the following Lemma, $Z - V$ has codimension ≥ 2 in Z :

Lemma 1.14. *The adjustable locus \mathcal{N}' is smooth over $\overline{\mathcal{M}}_{g,r+2}$ of relative dimension $3g - 3 + r + 1$ at points of \mathcal{N}' that are in the image of $Z - V$.*

Proof. Let us first consider the following result concerning \mathbf{P}^1 -bundles on \mathbf{P}^1 : Let k be an algebraically closed field. Let $X^{\log} \rightarrow \text{Spec}(k)$ (where we give $\text{Spec}(k)$ the trivial log structure) be the log-curve whose underlying scheme $X = \mathbf{P}_k^1$ and whose log structure is defined by a single marked point “ ∞ .” Let $T \rightarrow \text{Spec}(k)$ be a connected k -scheme. We shall denote pull-backs of objects over k to T by means of a subscript “ T .” Let (P, ∇_P) be a \mathbf{P}^1 -bundle with connection on X_T^{\log} equipped with a “Hodge” section $\alpha : X_T \rightarrow P$ whose Kodaira-Spencer morphism is an isomorphism except at the points “0” and “1” of $X_T = \mathbf{P}_T^1$, where the Kodaira-Spencer morphism has simple zeroes. Then we propose to show the following:

(*) The \mathbf{P}^1 -bundle with connection (P, ∇_P) is the pull-back to X_T^{\log} of a \mathbf{P}^1 -bundle with connection on X^{\log} . Moreover, the monodromy operator μ_{∞} of (P, ∇_P) at “ ∞ ” satisfies $\langle \mu_{\infty}, \mu_{\infty} \rangle \in \Gamma(T, \mathcal{O}_T^{\times})$.

We do this as follows: Let $\mathcal{T} \stackrel{\text{def}}{=} \Phi_{X_T}^* \tau_{X_T^{\log}/T}$. Note, moreover, that \mathcal{T} is a line bundle of relative degree p over T . The p -curvature of (P, ∇_P) is a morphism $\mathcal{P} : \mathcal{T} \rightarrow \text{Ad}(P)$. Let $\mathcal{L} \stackrel{\text{def}}{=} \alpha^* \tau_{P/X_T}$. By the assumption on the Kodaira-Spencer morphism, it follows that the line bundle \mathcal{L} has relative degree 3 over T . Moreover, α defines a Hodge filtration $F^{\bullet}(-)$ on $\text{Ad}(P)$ with subquotients \mathcal{L} , \mathcal{O}_{X_T} , and \mathcal{L}^{-1} . Since $p \geq 5$, it follows by degree considerations that any morphism between \mathcal{T} and any of these various subquotients is identically zero. Thus, \mathcal{P} is identically zero. Let us assume (by possibly replacing T by a finite flat double covering of T) that there exists a section ρ_{∞} of \mathcal{O}_T over T such that $2\rho_{\infty}^2 = \langle \mu_{\infty}, \mu_{\infty} \rangle$. Then it follows from the facts that

- (i) the Kodaira-Spencer morphism of α is nonzero at ∞ , and
- (ii) \mathcal{P} is identically zero (cf. the discussion preceding Proposition 1.3),

that $\rho_{\infty} \in \Gamma(T, \mathcal{O}_T^{\times})$. Thus, by *indigenizing* (P, ∇_P) (as in the discussion preceding Theorem 1.9), it follows that we are dealing with a family (parametrized by T) of dormant torally indigenous bundles on the 3-pointed smooth curve of genus 0 two of whose radii are 1, and whose third radius is an invertible function ρ_{∞} on T . But then it follows immediately from our computation of the p -curvature of a torally indigenous bundle at a marked point (the discussion preceding Proposition

1.3) that $\rho_\infty \in \mathbf{F}_p^\times$, and hence that (P, ∇_P) is the pull-back of a bundle with connection on X^{\log} (cf. Chapter I, Theorem 4.4). This completes the proof of (*).

Now we claim that (*) suffices to prove the Lemma. Indeed, to prove the Lemma, it suffices to consider a geometric fiber F of $\overline{\mathcal{S}}_{g,r+2}' \rightarrow \overline{\mathcal{M}}_{g,r+2}$ at a point ν of $\overline{\mathcal{M}}_{g,r+2}(k)$ which is in the image of $Z - V$. Then we must show that the adjustable locus $L \subseteq F$ (for the last marking section) is smooth over k of dimension $3g - 3 + r + 1$. But to say that the curve ν is in the image of $Z - V$ means that the two auxiliary sections of the curve ν lie on a copy of \mathbf{P}_k^1 (i.e., “ X ”) with no other marked points and only one node (“ ∞ ”). Let us think of the two auxiliary sections as the points “0” and “1” of $X = \mathbf{P}_k^1$. Then any torally indigenous bundle (say, over some connected base T) in the fiber L under consideration will have some restriction (K, ∇_K) to X_T , where ∇_K is a connection with logarithmic singularities at “0,” “1,” and “ ∞ .” Let us *assume* that

(†) The radius of the monodromy of (K, ∇_K) at ∞ is either 0 or an invertible function on T .

(For instance, this assumption is always satisfied if T is the spectrum of a field.) Then it follows that (K, ∇_K) is a *torally indigenous bundle* on X_T^{\log} . Moreover, (K, ∇_K) has radius 1 at 0 and 1. Note that (K, ∇_K) is completely symmetric with respect to the automorphism of X that transposes 0 and 1 and fixes ∞ . (This is because (K, ∇_K) is completely determined by its monodromy at 0, 1, and ∞ – cf. Chapter I, Theorem 4.4.) Thus, the fact that (K, ∇_K) is assumed to be adjustable at one of 0 and 1 implies that it is adjustable at *both* 0 and 1. Adjusting (K, ∇_K) at 0 and 1 thus gives us a (P, ∇_P) as in (*). Thus, (*) implies that the radius of (K, ∇_K) at ∞ is invertible. This in turn (applied to the case of T equal to the spectrum of a field) implies that *the assumption* (†) *is always satisfied*. Thus, we may apply (*) over an arbitrary connected T . Then (*) tells us that such a (K, ∇_K) or (P, ∇_P) has no moduli, hence that the locus L is scheme-theoretically defined by the condition that the restriction (K, ∇_K) be equal to one of a finite number of possibilities. This condition clearly defines a smooth (even linear) subvariety of F of the correct dimension. This completes the proof of the Lemma. \circ

Let $(Q', \nabla_{Q'})$ be the adjustment of $(Q, \nabla_Q)|_{\mathcal{N}'}$ at the last marking section. Thus, $(Q', \nabla_{Q'})$ is a \mathbf{P}^1 -bundle with connection on $(Y')^{\log}$. By restricting to V , we thus obtain a \mathbf{P}^1 -bundle with connection on $Z^{\log}|_V$. Since $Z - V$ has codimension ≥ 2 in Z and Z is regular (by the preceding Lemma) at the points of $Z - V$, we can thus push-forward this \mathbf{P}^1 -bundle with connection on $Z^{\log}|_V$ to obtain a \mathbf{P}^1 -bundle with connection (R, ∇_R) on Z^{\log} . We shall refer to (R, ∇_R) as the *contracted adjustment* of $(Q, \nabla_Q)|_{\mathcal{N}'}$.

Note that (R, ∇_R) has radius 1 at σ_Z , and radii ρ at the first r marking sections of Z^{\log} . Let $\mathcal{N}'' \subseteq \mathcal{N}'$ be the adjustable locus of (R, ∇_R)

for σ_Z . Endow \mathcal{N}'' with the pull-back to \mathcal{N}'' of the log structure of \mathcal{N}' . Let $f^{\log} : X^{\log} \rightarrow (\mathcal{N}'')^{\log}$ be the r -pointed semi-stable curve of genus g obtained from $Z^{\log}|_{\mathcal{N}''}$ by forgetting the section $\sigma_Z|_{\mathcal{N}''}$. Let (H, ∇_H) be the adjustment of $(R, \nabla_R)|_{\mathcal{N}''}$ at $\sigma_Z|_{\mathcal{N}''}$. Thus, (H, ∇_H) is a \mathbf{P}^1 -bundle with connection on X^{\log} . Let us consider the Verschiebung (i.e., the determinant of the p -curvature) of (H, ∇_H) . This Verschiebung forms a section of the vector bundle $\mathcal{Q} \stackrel{\text{def}}{=} \Phi_{\mathcal{N}''}^* f_* \omega_{X^{\log}/(\mathcal{N}'')^{\log}}^{\otimes 2}(-M_f)$ over \mathcal{N}'' . Let $\mathcal{N}_{g,r}^{\rho, \text{psd}} \subseteq \mathcal{N}''$ be the zero locus of the Verschiebung. Now we have the following

Lemma 1.15. *The morphism $\mathcal{N}_{g,r}^{\rho, \text{psd}} \rightarrow \overline{\mathcal{M}}_{g,r+2}$ is quasi-finite.*

Proof. At any rate, $\mathcal{N}_{g,r}^{\rho, \text{psd}} \subseteq S = \overline{\mathcal{S}}_{g,r+2}^{\rho'}$. Moreover, for any $\nu \in \mathcal{N}_{g,r}^{\rho, \text{psd}}(k)$ (where k is an algebraically closed field), we claim that the original torally indigenous bundle (Q, ∇_Q) is nilpotent at ν . Note that by Proposition 1.7, this will complete the proof of the Lemma. On the other hand, to check that $(Q, \nabla_Q)_{\nu}$ is nilpotent, it suffices to check that its p -curvature is nilpotent over an dense open subset of every irreducible component of Y_{ν} . But the only irreducible component of Y_{ν} that does not map finitely to an irreducible component of X_{ν} (where we know that the p -curvature is nilpotent) is the sort of irreducible component examined in the proof of Lemma 1.14. On the other hand, in the proof of Lemma 1.14, we already saw that (since $\nu \in \mathcal{N}'(k)$), $(Q, \nabla_Q)_{\nu}$ has zero p -curvature on that sort of irreducible component. \circ

On the other hand, note that $\mathcal{N}_{g,r}^{\rho, \text{psd}}$ as a closed substack of S is locally defined by precisely $3g - 3 + r + 2$ equations. It thus follows that we have the following result:

Theorem 1.16. *Suppose that $p \geq 5$, and $2g - 2 + r \geq 1$. Then the morphism $\mathcal{N}_{g,r}^{\rho, \text{psd}} \rightarrow \overline{\mathcal{M}}_{g,r+2}$ is finite and flat, and the degree of $(\mathcal{N}_{g,r}^{\rho, \text{psd}})_{\text{red}}$ over $\overline{\mathcal{M}}_{g,r+2}$ is $\leq p^{3g-3+r+2}$. Moreover, if $p > 2^{3g-3+r+2}$, then every irreducible component of $(\mathcal{N}_{g,r}^{\rho, \text{psd}})_{\text{red}}$ is generically étale over $\overline{\mathcal{M}}_{g,r+2}$.*

Proof. Note that $\mathcal{N}_{g,r}^{\rho, \text{psd}}$ is a closed substack of S which (as least point-theoretically – as in the proof of Lemma 1.15) is contained in $\overline{\mathcal{N}}_{g,r+2}^{\rho'}$. Thus, $\mathcal{N}_{g,r}^{\rho, \text{psd}}$ is proper over \mathbf{F}_p . This shows that the morphism in question is finite, and that its degree is bounded as stated. The last statement follows from Proposition 1.8 and the fact that $(\mathcal{N}_{g,r}^{\rho, \text{psd}})_{\text{red}}$ is contained in $\overline{\mathcal{N}}_{g,r+2}^{\rho'}$. \circ

Definition 1.17. We shall refer to (H, ∇_H) as the *tautological PTCS* (or, *pseudotorally crys-stable*) bundle of radii ρ and level $\chi - 1$ on X^{\log} .

Moreover, let $(\mathcal{N}_{g,r}^{\rho, \text{psd}})' \subseteq \mathcal{N}_{g,r}^{\rho, \text{psd}}$ be the open (dense, by Theorem 1.16) substack which is the inverse image of $\mathcal{M}_{g,r+2}$ via the above

morphism. Let $\mathcal{N}_{g,r}^{\rho,\text{crs}}$ be the stack of nilpotent torally crys-stable bundles of radii ρ and level $\chi - 1$ on *smooth* r -pointed curves of genus g whose Kodaira-Spencer loci are étale (divisors of degree 2) over the base. Then note that $(H, \nabla_H)|_{(\mathcal{N}_{g,r}^{\rho,\text{psd}})'}$ defines a morphism

$$(\mathcal{N}_{g,r}^{\rho,\text{psd}})' \rightarrow \mathcal{N}_{g,r}^{\rho,\text{crs}}$$

Proposition 1.18. *This morphism $(\mathcal{N}_{g,r}^{\rho,\text{psd}})' \rightarrow \mathcal{N}_{g,r}^{\rho,\text{crs}}$ is finite étale of degree 2.*

Proof. This follows immediately by sorting through all the definitions. The reason the morphism is of degree two is because over $(\mathcal{N}_{g,r}^{\rho,\text{psd}})'$, the two points that make up the Kodaira-Spencer locus are ordered. \circ

§2. Nilpotent Connections of Higher Order

In this §, we introduce a number of *ad hoc* notions (connections of higher order, p^n -curvature, etc.) for the purpose of showing that a certain moduli functor of connections of higher order is *smooth*. The point of doing this is that it is much easier to construct such connections on singular curves or smooth curves with (toral) marked points than on smooth curves without marked points, even though we are ultimately really interested (for instance, in the case of Lubin-Tate uniformizations) only in such objects on smooth curves without marked points. It will turn out that all the notions that we introduce and all the smoothness results that we prove will be completely trivial for smooth curves without marked points; thus, what we do in this § will only be nontrivial for singular curves or smooth curves with (toral) marked points. In the case of smooth curves without marked points, however, the moduli functor of connections of higher order that we introduce in this § is somewhat unnatural relative to the uniformization theory that is our ultimate goal. Although, in some sense, this functor carries roughly the same information as is contained in (the appropriate portion of) the moduli stacks associated to VF-patterns that we introduce in Chapter III, the stacks associated to VF-patterns are much more natural from the point of view of uniformization theory. On the other hand, “historically,” the author first discovered the notions introduced here, and only later did he realize that the point of view of Chapter III was more natural from the point of view of uniformization theory.

§2.1. Higher Order Connections

Let n be a nonnegative integer. Let S^{\log} be a fine log scheme whose underlying scheme is noetherian and flat over $\mathbf{Z}/p^{n+1}\mathbf{Z}$. If $0 \leq m \leq n$, then we shall denote by a subscript m the result of reducing an object over $\mathbf{Z}/p^{n+1}\mathbf{Z}$ modulo p^{m+1} . Let $X_0^{\log} \rightarrow S_0^{\log}$ be *locally stable of dimension one*. (By this, we mean that $X_0^{\log} \rightarrow S_0^{\log}$ is étale locally isomorphic to an open sub-(log)-scheme of a stable log-curve (Chapter I, Definition 1.1). See, e.g., [Mzk1], Chapter I, Definition 1.1, for more details.). Suppose that $X_0^{\log} \rightarrow S_0^{\log}$ has r marked points. Let us denote by $\Phi_X : X_0 \rightarrow X_0$ (respectively, $\Phi_S : S_0 \rightarrow S_0$) the absolute Frobenius on X_0 (respectively, S_0). Let us denote by $X_0^{F^m} \rightarrow S_0$ the result of pulling back $X_0 \rightarrow S_0$ by Φ_S^m . Let us denote by $\Phi_{X/S}^m : X_0 \rightarrow X_0^{F^m}$ the relative Frobenius morphism of X_0 over S_0 obtained by raising sections of $\mathcal{O}_{X_0^{F^m}}$ to the power p^m .

Let us fix a set of invertible radii $\rho_1, \dots, \rho_r \in \mathcal{O}_S^\times$. We will denote the set of ρ_i 's by ρ . Let (P_0, ∇_{P_0}) be a \mathbf{P}^1 -bundle with connection on X_0^{\log} .

Definition 2.1. We shall call a crystal (P, ∇_P) (in \mathbf{P}^1 -bundles) on $\text{Crys}(X_0^{\log}/S^{\log})$ a *pre- n -connection* on (P_0, ∇_{P_0}) (of radii ρ) if $(P, \nabla_P) \otimes \mathbf{F}_p = (P_0, \nabla_{P_0})$ and the monodromy operator μ_i of (P, ∇_P) at the i^{th} marked point satisfies $\langle \mu_i, \mu_i \rangle = 2\rho_i^2$.

Suppose that we have a pre- n -connection (P, ∇_P) on (P_0, ∇_{P_0}) (of radii ρ). Let $U_0 \subseteq X_0$ be the open subset which is the complement of the marked points and nodes. Let $U^{\log} \rightarrow S^{\log}$ be a $\mathbf{Z}/p^{n+1}\mathbf{Z}$ -flat thickening of $U_0^{\log} \rightarrow S^{\log}$. Let $\mathcal{A}^{[0]} = \text{Ad}(P_0)|_{U_0}$. Thus, $\mathcal{A}^{[0]}$ is equipped with a connection $\nabla_{\mathcal{A}^{[0]}}$. We would like to define, *inductively on m* (a nonnegative integer $\leq n$) the following objects:

- (1) $_m$ a subsheaf (in the Zariski topology) $\mathcal{A}^{[m]} \subseteq \mathcal{A}^{[0]}$ which forms a coherent sheaf of $\mathcal{O}_{U_0^{F^m}}$ -modules;
- (2) $_m$ a connection $\nabla_{\mathcal{A}^{[m]}}$ on $\mathcal{A}^{[m]}$ (regarded as a coherent sheaf of $\mathcal{O}_{U_0^{F^m}}$ -modules)

subject to the “inductive condition” that the p -curvature \mathcal{P}_m of the connection $\nabla_{\mathcal{A}^{[m-1]}}$ (regarded as a connection on a bundle of *Lie algebras*) be identically zero. (We regard this condition as vacuous when $m = 0$.) Moreover, assuming this “inductive condition,” we shall prove (inductively, for $m \geq 1$) that Zariski locally on U , there exists a basis of horizontal sections of $\mathcal{A}^{[m-1]}$ that lift to horizontal sections of $\text{Ad}(P_{m-1})|_U$. We start the induction by considering the case $m = 0$. Then we have already defined $\mathcal{A}^{[0]}$ and $\nabla_{\mathcal{A}^{[0]}}$.

Now suppose that the case $m - 1$ has been dealt with. Thus, in particular, the pair $(\mathcal{A}^{[m-1]}, \nabla_{\mathcal{A}^{[m-1]}})$ has been defined and has p -curvature 0. Let $\mathcal{A}^{[m]} \subseteq \mathcal{A}^{[m-1]}$ be the subsheaf of horizontal sections.

Thus, $\mathcal{A}^{[m]}$ is a coherent sheaf of $\mathcal{O}_{X_0^{F^m}}$ -modules. Moreover, by the induction hypothesis, we know that $\mathcal{A}^{[m]}$ admits (locally on U_0) a basis of sections that lift to horizontal sections of $\mathrm{Ad}(P_{m-1})$. By taking linear combinations of these sections with coefficients of the form ϕ^{p^m} (where ϕ is a section of \mathcal{O}_U), we see that in fact, every section of $\mathcal{A}^{[m]}$ can (Zariski locally) be lifted to a horizontal section of $\mathrm{Ad}(P_{m-1})$.

Now we would like to define a connection on $\mathcal{A}^{[m]}$ as follows. Let α_0 be a local section of $\mathcal{A}^{[m]}$ over some open set $V_0 \subseteq U_0$. Then V_0 defines an open subset $V \subseteq U$. Lift α_0 to a section α of $\mathrm{Ad}(P_m)$ over V which is horizontal modulo p^m . Applying the connection on $\mathrm{Ad}(P)$ to α gives us a section α' of $p^m \cdot \mathrm{Ad}(P_m) \otimes_{\mathcal{O}_U} \omega_{U/S}$ over V . Thus, by dividing by p^m , we see that α' gives us a section β of $\mathcal{A}^{[0]} \otimes_{\mathcal{O}_{U_0}} \omega_{U_0/S_0}$ over V . On the other hand, note that $\mathcal{A}^{[0]} = (\Phi_{X/S}^m)^* \mathcal{A}^{[m]}$. Thus, we may apply the *Cartier isomorphism* to β a total of m times. This gives us a section γ of $\mathcal{A}^{[m]} \otimes_{\mathcal{O}_{U_0^{F^m}}} \omega_{U_0^{F^m}/S_0}$ over V . It is easy to see that the assignment $\alpha_0 \mapsto \gamma$ defines a connection on $\mathcal{A}^{[m]}$ which is independent of the choice of α , $U^{\log} \rightarrow S^{\log}$. (Indeed, if ϵ is a section of $\mathrm{Ad}(P_m)$ which is horizontal modulo p^m , and ϕ is a section of \mathcal{O}_U , then it suffices to observe that sections α of $\mathrm{Ad}(P_m)$ of the form $p^{m-a} \cdot \phi^{p^a} \cdot \epsilon$ (where $0 \leq a < m$) give rise to $\gamma = 0$.) Thus, we obtain a connection $\nabla_{\mathcal{A}^{[m]}}$ on $\mathcal{A}^{[m]}$. The p -curvature of this connection is a horizontal section

$$\mathcal{P}_{m+1}^U \in \Gamma(U_0, \mathcal{A}^{[m]} \otimes \Phi_X^*(\omega_{X_0^{\log}/S_0^{\log}})^{F^m}|_U)$$

We shall call \mathcal{P}_{m+1}^U the p^{m+1} -curvature (over U_0) of the pre- n -connection (P, ∇_P) on (P_0, ∇_{P_0}) .

Suppose that \mathcal{P}_{m+1}^U is identically zero. Let α_0 be a local horizontal section of $\mathcal{A}^{[m]}$ over some open set $V_0 \subseteq U_0$. Lift α_0 to a section α of $\mathrm{Ad}(P_m)$ over V which is horizontal modulo p^m . Applying the connection on $\mathrm{Ad}(P)$ to α gives us a section α' of $p^m \cdot \mathrm{Ad}(P_m) \otimes_{\mathcal{O}_U} \omega_{U/S}$ over V with the property that when we divide α' by p^m and apply the Cartier isomorphism m times, we get zero. This means that to integrate $\frac{1}{p^m} \alpha'$ the largest denominator that we need is p^{m-1} . In other words, α' can be written in the form $\nabla(p \cdot \delta)$, where ∇ is the connection on $\mathrm{Ad}(P)$, and δ is some section of $\mathrm{Ad}(P_m)$ over V . Thus, it follows that $\alpha - p\delta$ (which is still a lifting of α_0) is a horizontal section of $\mathrm{Ad}(P_m)$. This completes our treatment of the case m .

We are now ready to make the main definition of this §:

Definition 2.2. A pre- n -connection (P, ∇_P) on (P_0, ∇_{P_0}) is called an n -connection on (P_0, ∇_{P_0}) if it satisfies the above inductive condition – namely, that the p^m -curvature be identically zero for all m with $1 \leq m \leq n$. We shall call an n -connection *dormant* (respectively, *nilpotent*) if its p^{n+1} -curvature (over U_0) vanishes (respectively, has vanishing determinant).

In particular, to give a 0-connection on (P_0, ∇_{P_0}) is to simply give (P_0, ∇_{P_0}) itself.

Note, moreover, that for any n -connection (P, ∇_P) on (P_0, ∇_{P_0}) , we obtain (by the above discussion) a horizontal section

$$\mathcal{P}_{n+1}^U \in \Gamma(U_0, \mathcal{A}^{[n]} \otimes_{\mathcal{O}_{X_0^{F^n}}} \Phi_X^*(\omega_{X_0^{\log}/S_0^{\log}})^{F^n}|_U)$$

which we call the p^{n+1} -curvature (over U_0) of the n -connection (P, ∇_P) on (P_0, ∇_{P_0}) . Now note that we can regard the sheaf $\mathcal{A}^{[n]} \otimes \Phi_X^*(\omega_{X_0^{\log}/S_0^{\log}})^{F^n}|_U$ as a subsheaf of its pull-back by $\Phi_{X/S}^n$. But this pull-back is none other than $\mathcal{A}^{[0]} \otimes (\Phi_{X/S}^*)^n(\Phi_X^*(\omega_{X_0^{\log}/S_0^{\log}})^{F^n})|_U = (\text{Ad}(P_0) \otimes (\Phi_X^*)^{n+1}(\omega_{X_0^{\log}/S_0^{\log}}))|_U$. By abuse of notation, we shall often regard \mathcal{P}_{n+1}^U as a section of $(\text{Ad}(P_0) \otimes (\Phi_X^*)^{n+1}(\omega_{X_0^{\log}/S_0^{\log}}))|_U$.

Now let us review what we have done so far. We started with a $\mathbf{Z}/p^{n+1}\mathbf{Z}$ -flat fine log scheme S^{\log} ; a locally stable morphism of dimension one $X_0^{\log} \rightarrow S_0^{\log}$; and a \mathbf{P}^1 -bundle with connection (P_0, ∇_{P_0}) on X_0^{\log} . Relative to this data, we defined the notion of an n -connection (P, ∇_P) of radii ρ on (P_0, ∇_{P_0}) . Before proceeding, we would like to consider what conditions are imposed on the radii ρ by the fact that (P, ∇_P) defines an n -connection on (P_0, ∇_{P_0}) .

Lemma 2.3. *Suppose that S is connected. Then:*

- (1) *If either (i) the p^{n+1} -curvature \mathcal{P}_{n+1} of (P, ∇_P) is identically zero, or (ii) the p^{n+1} -curvature \mathcal{P}_{n+1} of (P, ∇_P) is square nilpotent and S is reduced, then the radii ρ_i all lie in $(\mathbf{Z}/p^{n+1}\mathbf{Z})^\times$.*
- (2) *Suppose that \mathcal{P}_{n+1} is identically zero. Then locally, near a marked point of radius ρ , (P, ∇_P) is isomorphic to the projective bundle associated to $(t^{-a} \cdot \mathcal{O}_X) \oplus (t^{-b} \cdot \mathcal{O}_X)$ (equipped with the connection induced by the trivial connection on $\mathcal{O}_X \oplus \mathcal{O}_X$). Here, $X^{\log} \rightarrow S^{\log}$ is a flat lifting of $X_0^{\log} \rightarrow S_0^{\log}$; t is a local parameter on X whose zero locus is the marked point in question; and a and b are integers equal to $\pm \rho \pmod{p^{n+1}}$ such that $0 \leq a \leq p^{n+1} - 1$, $a + b = p^{n+1}$.*

Finally, if \mathcal{P}_{n+1} is identically zero, and, moreover, we are given an equisingular lifting $X^{\log} \rightarrow S^{\log}$ (which allows us to regard the crystal (P, ∇_P) as a crys-stable bundle on $X^{\log} \rightarrow S^{\log}$), then the monodromy of (P, ∇_P) at a node has radius $\in (\mathbf{Z}/p^{n+1}\mathbf{Z})^\times$.

Proof. Clearly, the final statement of the Lemma follows immediately from (1), so we will concentrate on (1) and (2). First, observe that since S is connected, it suffices to prove the Lemma in the case that S is local (i.e., equal to Spec of a local ring). Thus, let us assume that S is local. Let us fix a lifting $X^{\log} \rightarrow S^{\log}$ of $X_0^{\log} \rightarrow S_0^{\log}$ (so $X \rightarrow S$ is flat). Let us consider what happens at a marked point $x \in X(S)$. By replacing X by an open neighborhood of x , we may assume that X is

affine, has no nodes, and has no marked points other than x . (Thus, for the rest of the proof, $r = 1$, $\rho_1 = \rho$.) By abuse of notation, we shall denote by (P, ∇_P) the \mathbf{P}^1 -bundle with connection defined by the crystal (P, ∇_P) on X^{\log} . We may assume (by further shrinking X) that (P, ∇_P) is the projectivization of a rank two bundle with connection $(\mathcal{E}, \nabla_{\mathcal{E}})$ on X whose determinant is trivial. Let t be a local parameter on X which is zero at x .

Now, we would like to prove (1) and (2) by induction on n . Let us begin with the case $n = 0$. In this case, (1) follows from the discussion preceding Proposition 1.3. Thus, we proceed to prove (2). Now for any integer i , note that (since x is a marked point of X^{\log}), the connection $\nabla_{\mathcal{E}}$ induces a connection $\nabla_{\mathcal{E}}^{[i]}$ on $t^i \cdot \mathcal{E}$. Suppose that a and b are nonnegative integers $\leq p-1$ such that $a+b = p$ and a is congruent modulo p to the radius at the marked point. Without loss of generality, we shall also assume that $a < b$. Then I *claim* that there exists a section $\alpha \in t^a \cdot \mathcal{E}$ with the following properties: (i) the zero locus of α (as a section of $t^a \cdot \mathcal{E}$) avoids x ; (ii) on the closed subscheme $V(t^{b-a}) \subseteq X$, α is horizontal with respect to $\nabla_{\mathcal{E}}^{[a]}$. Indeed, since the monodromy operator of $\nabla_{\mathcal{E}}^{[a]}$ at x has eigenvalues 0 and $a-b$ (modulo p), it is clear how to construct “ $\alpha|_x$.” Once one has constructed $\alpha|_{V(t^i)}$ (for $i < b-a$), the fact that one has a lifting $\alpha|_{V(t^{i+1})}$ with the appropriate properties follows from the fact that the monodromy operator of $\nabla_{\mathcal{E}}^{[a+i]}$ has eigenvalues i and $a-b+i$, neither of which is $\equiv 0 \pmod{p}$. This completes the proof of the claim. Now let $\mathcal{G} \subseteq t^a \cdot \mathcal{E}$ be the \mathcal{O}_X -submodule generated by α and $t^b \cdot \mathcal{E}$. Thus, \mathcal{G} is a rank two vector bundle equipped with a logarithmic connection $\nabla_{\mathcal{G}}$ (induced by $\nabla_{\mathcal{E}}$). On the other hand, by construction, the monodromy operator of $\nabla_{\mathcal{G}}$ at x is *square nilpotent* (cf. the discussion preceding Definition 1.13). Since the p -curvature of $(\mathcal{G}, \nabla_{\mathcal{G}})$ is assumed to be zero, however, we conclude (cf. the discussion preceding Proposition 1.3) that the monodromy operator is, in fact, *zero*. Thus, \mathcal{G} admits a horizontal basis γ_1, γ_2 (in a neighborhood of x) such that γ_1 coincides with α over $V(t^{b-a})$ and $\gamma_2 \in t^b \cdot \mathcal{E}$. This basis then defines a local isomorphism of $(\mathcal{E}, \nabla_{\mathcal{E}})$ with $(t^{-a} \cdot \mathcal{O}_X) \oplus (t^{-b} \cdot \mathcal{O}_X)$, as desired. This completes the proof of the Lemma in the case $n = 0$.

Now we assume that (1) and (2) are known for $n-1$ (where $n > 0$); we would then like to conclude (1) and (2) for n . Again, we use the notation $\nabla_{\mathcal{E}}^{[i]}$ for the connection on $t^i \cdot \mathcal{E}$ induced by $\nabla_{\mathcal{E}}$. By the induction hypothesis, there exists a nonnegative integer $a \leq p^n - 1$ (such that $a < b \stackrel{\text{def}}{=} p^n - a$) and a section $\alpha \in t^a \cdot \mathcal{E}$ such that $\mathcal{H} \stackrel{\text{def}}{=} (t^a \cdot \mathcal{E})/(\mathcal{O}_X \cdot \alpha)$ is a line bundle on X , and α is horizontal modulo p^n . I *claim* that there exists a section $\alpha' \in t^a \cdot \mathcal{E}$ with the following properties: (i) $\alpha \equiv \alpha' \pmod{p}$; (ii) α' is horizontal modulo p^n ; (iii) the restriction of $\nabla_{\mathcal{E}}^{[a]}(\alpha')$ to $V(t^{b-a})$ is contained in $\mathcal{O}_X \cdot \alpha$. Indeed, let $\beta \stackrel{\text{def}}{=} \nabla_{\mathcal{E}}^{[a]}(\alpha) \in p^n \cdot (t^a \cdot \mathcal{E}_0) \otimes \omega_{X_0^{\log}/S_0^{\log}}$. Now by the induction hypothesis, $(t^a \cdot \mathcal{E})_{n-1}$ is horizontally isomorphic to $(\mathcal{O}_{X_{n-1}}) \oplus (t^{a-b} \cdot \mathcal{O}_{X_{n-1}})$ (where the first summand is generated by α). Thus, to integrate the image of β in $\mathcal{H} \otimes \omega_{X^{\log}/S^{\log}}$ up to order $b-a$, it suffices to introduce denominators dividing p^{n-1} (since $b-a < p^n$). On

the other hand, since β is divisible by p^n to begin with, we conclude that there exists a $\gamma \in p \cdot t^a \cdot \mathcal{E}$ which is horizontal modulo p^n and such that the image of $\nabla_{\mathcal{E}}^{[a]}(\gamma)$ in $\mathcal{H} \otimes \omega_{X^{\log}/S^{\log}}$ coincides with the image of β up to order $b - a$. Taking $\alpha' \stackrel{\text{def}}{=} \alpha - \gamma$ thus proves the claim. In the following, we replace α by α' , and thus, assume that α has the properties assumed of α' .

Let $\mathcal{G} \subseteq t^a \cdot \mathcal{E}$ be the subbundle generated by α and $t^b \cdot \mathcal{E}$. Let $\nabla_{\mathcal{G}}$ be the connection induced on \mathcal{G} . The image of $t^b \cdot \mathcal{E}$ defines a subspace on $\mathcal{K} \subseteq \mathcal{G}|_x$ which is preserved by the monodromy operator of $\nabla_{\mathcal{G}}$ at x . Thus, relative to the filtration defined by this subspace the monodromy operator of $\nabla_{\mathcal{G}}$ at x has the following form:

$$\begin{pmatrix} b + \rho & \zeta \\ 0 & a - \rho \end{pmatrix}$$

Here $a - \rho$, $b + \rho$, and ζ are $\equiv 0$ modulo p^n . Let $\mathcal{G}^{[n]} \subseteq \mathcal{G}_0$ be the subsheaf of sections that lift to horizontal sections modulo p^n . Thus, $\text{Ad}(\mathcal{G}^{[n]})$ may be identified with $\mathcal{A}^{[n]}$. Then note that the same recipe used to define $\nabla_{\mathcal{A}^{[n]}}$ allows us to construct a (logarithmic) connection $\nabla_{\mathcal{G}^{[n]}}$ on the $\mathcal{O}_{X_0^{F^n}}$ -module $\mathcal{G}^{[n]}$. Write $b + \rho = p^n \cdot \xi_1$; $a - \rho = p^n \cdot \xi_2$; $\zeta = p^n \cdot \xi_3$, where $\xi_1, \xi_2, \xi_3 \in \mathcal{O}_{S_0}$. Thus, in particular, the fact that $a + b = p^n$ implies that $\xi_1 + \xi_2 = 1$. Moreover, the monodromy operator of $\nabla_{\mathcal{G}^{[n]}}$ at x is easily computed to be:

$$\begin{pmatrix} \xi_1 & \xi_3 \\ 0 & \xi_2 \end{pmatrix}$$

Since we assume that $\text{Ad}(\mathcal{G}^{[n]}, \nabla_{\mathcal{G}^{[n]}})$ either (i) has p -curvature zero or (ii) has square nilpotent p -curvature and S is reduced, we conclude from the discussion preceding Proposition 1.3 that $\xi_1, \xi_2 \in \mathbb{F}_p$. This implies that $\rho \in \mathbb{Z}/p^{n+1}\mathbb{Z}$, as desired. This completes the proof of (1) (for the case n).

Finally, we would like to verify (2) (for the case n). Let a' be the nonnegative integer $\leq p - 1$ which is $\equiv -\xi_2$ modulo p . Now observe that by the argument used to verify (2) in the case $n = 0$, it follows that there exists a horizontal section $\epsilon \in t^{p^n \cdot a'} \cdot \mathcal{G}^{[n]}$ with the following property: when ϵ is regarded as a section of $t^{p^n \cdot a'} \cdot \mathcal{G}_0$, the restriction of ϵ to x does not lie in the subspace $t^{p^n \cdot a'} \cdot \mathcal{K}_0 \subseteq t^{p^n \cdot a'} \cdot \mathcal{G}_0$ at any point of S . Next, observe that it follows from the definition of $\nabla_{\mathcal{G}^{[n]}}$ that ϵ lifts to a section $\theta \in t^{p^n \cdot a'} \cdot \mathcal{G}$ which is horizontal modulo p^n and, moreover, has the property that $\frac{1}{p^n} \cdot \nabla_{\mathcal{G}}^{[p^n \cdot a']}(\theta)$ is annihilated by the n^{th} power of the Cartier operator. It thus follows (by an argument similar to the argument used to construct α' above) that there exists a horizontal section $\theta' \in t^{p^n \cdot a'} \cdot \mathcal{G}$ which is $\equiv \theta \pmod{p}$. In other words, if $a'' \stackrel{\text{def}}{=} a + p^n \cdot a'$, then θ' defines a locally split horizontal injection $(t^{-a''} \cdot \mathcal{O}_X) \hookrightarrow \mathcal{E}$. By an analogous argument, one can construct a locally split horizontal

injection $(t^{a''-p^{n+1}} \cdot \mathcal{O}_X) \hookrightarrow \mathcal{E}$. This completes the proof of (2) (for the case n), and hence of the entire Lemma. \square

§2.2. De Rham Cohomology Computations

We maintain the notations of the preceding paragraph. Moreover, we assume in addition, henceforth, that *all the radii ρ_i belong to $(\mathbf{Z}/p^{n+1}\mathbf{Z})^\times$* . In particular, we suppose that we are given an n -connection (P, ∇_P) on (P_0, ∇_{P_0}) . Let us lift $X_0^{\log} \rightarrow S_0^{\log}$ to a locally stable morphism of dimension one $X^{\log} \rightarrow S^{\log}$, so that the crystal (P, ∇_P) becomes a \mathbf{P}^1 -bundle with connection on X^{\log} (which, by abuse of notation, we also denote by (P, ∇_P)). We would like to consider the local de Rham cohomology of (P, ∇_P) . For each marked point $x_i : S \rightarrow X$, let $\mu_i \in \text{Ad}(P)|_{x_i}$ be the monodromy operator of (P, ∇_P) at x_i . Let $\text{Ad}^q(P) \subseteq \text{Ad}(P)$ be the subsheaf of sections s of $\text{Ad}(P)$ such that for each marked point x_i of X^{\log} , $\langle s|_{x_i}, \mu_i \rangle = 0$. Then recall that, in addition to the *naive de Rham complex*:

$$d^{\text{nv}} : \text{Ad}(P) \rightarrow \text{Ad}(P) \otimes_{\mathcal{O}_X} \omega_{X^{\log}/S^{\log}}$$

we also have a *toral de Rham complex*:

$$d^{\text{tr}} : \text{Ad}(P) \rightarrow \text{Ad}^q(P) \otimes_{\mathcal{O}_X} \omega_{X^{\log}/S^{\log}}$$

These complexes of sheaves on X have filtrations given by multiplying by powers of p . By considering the spectral sequences of each of these filtered complexes of sheaves, we obtain various differentials: that is, we have (say, in the naive case)

$$d^{\text{nv},[0]} \stackrel{\text{def}}{=} d_0^{\text{nv}} : \text{Ad}(P_0) \rightarrow \text{Ad}(P_0) \otimes_{\mathcal{O}_{X_0}} \omega_{X_0^{\log}/S_0^{\log}}$$

whose kernel and cokernel we denote by $K^{\text{nv},[1]}$ and $C^{\text{nv},[1]}$, respectively. Then, for each integer i such that $0 \leq i \leq n$, we have a differential

$$d^{\text{nv},[i]} : K^{\text{nv},[i]} \rightarrow C^{\text{nv},[i]}$$

whose kernel and cokernel we denote by $K^{\text{nv},[i+1]}$ and $C^{\text{nv},[i+1]}$. Roughly speaking, $d^{\text{nv},[i]}$ is the map obtained by applying d^{nv} to sections of $\text{Ad}(P_0)$ that lift to sections of $\text{Ad}(P)$ which are horizontal modulo p^i , then dividing by p^i , and finally reducing modulo p . Similarly, we have $d^{\text{tr},[i]}$, $K^{\text{tr},[i+1]}$, and $C^{\text{tr},[i+1]}$ in the toral case.

Now we have the following key result:

Lemma 2.4. *Suppose that $S_0 = \text{Spec}(k)$, where k is a perfect field, and that the p^{n+1} -curvature \mathcal{P}_{n+1} of (P, ∇_P) is identically zero. Suppose also that the deformation X^{\log} of X_0^{\log} is equisingular, in the sense that the nodes of X are $\mathbf{Z}/p^{n+1}\mathbf{Z}$ -flat. Then we have natural isomorphisms of sheaves on X_0 :*

$$C^{\text{nv},[n+1]} \cong (\text{Ad}(P_0) \otimes_{\mathcal{O}_{X_0}} (\Phi_X^{n+1})^* \omega_{X_0^{\log}/S_0^{\log}})^{\text{hor}}$$

and

$$C^{\text{tr},[n+1]} \cong (\text{Ad}^q(P_0) \otimes_{\mathcal{O}_{X_0}} (\Phi_X^{n+1})^* \omega_{X_0^{\log}/S_0^{\log}})^{\text{hor}}$$

where the superscript “hor” denotes the subsheaf of sections whose restriction to $U_0 \subseteq X_0$ (the complement of the nodes and marked points) comes from a section of $\mathcal{A}^{[n+1]} \otimes_{\mathcal{O}_{X_0^{F^{n+1}}}} (\omega_{X_0^{\log}/S_0^{\log}})^{F^{n+1}}|_U$ (in the notation of preceding Definition 2.2).

Proof. We shall give the proof in the naive case; the toral case is entirely similar. Also, since it will be clear that the isomorphism will be natural (at least at the smooth, unmarked points of X), we can (by descent) assume that k is algebraically closed. Let $U \subseteq X$ be the open subscheme of smooth, unmarked points of X . Over U , it is clear from the discussion preceding Definition 2.2 what happens: Namely, the differential $d^{\text{nv},[i]}$ may be identified with the connection

$$\nabla_{\mathcal{A}^{[i]}} : \mathcal{A}^{[i]} \rightarrow \mathcal{A}^{[i]} \otimes_{\mathcal{O}_{X_0^{F^i}}} (\omega_{X_0^{\log}/S_0^{\log}})^{F^i}|_{U_0}$$

on $\mathcal{A}^{[i]}$. This proves the Lemma over U .

It remains to extend the isomorphism over all of X . Let us first consider what happens at a marked point $x \in X(S)$. Let V be an affine neighborhood of x such that $V - \{x\} \subseteq U$ and $(P, \nabla_P)|_V$ is the projectivization of a rank two bundle with connection $(\mathcal{E}, \nabla_{\mathcal{E}})$ on V . Let t be a local parameter on V which is zero at x . By Lemma 2.3, (2), one knows that one can write $(\mathcal{E}, \nabla_{\mathcal{E}})$ in the following form: $(t^a \cdot \mathcal{O}_V) \oplus (t^b \cdot \mathcal{O}_V)$ (where a, b are integers such that $0 \leq b \leq a \leq p^{n+1} - 1$), equipped with the connection induced by the trivial connection on $\mathcal{O}_V \oplus \mathcal{O}_V$. Thus, $\text{Ad}(P)|_V = \text{Ad}(\mathcal{E})$ takes on the form $(t^c \cdot \mathcal{O}_V) \oplus \mathcal{O}_V \oplus (t^{-c} \cdot \mathcal{O}_V)$, where $c = a - b$, and $2\rho_j \equiv \pm c \pmod{p^{n+1}}$ (if x is the j^{th} marked point). In particular, c is nonzero, so $1 \leq c \leq p^{n+1} - 1$. Thus, we compute explicitly that the $C^{\text{nv},[n+1]}$ is

$$\{(t \cdot \mathcal{O}_{V_0}) \oplus \mathcal{O}_{V_0} \oplus \mathcal{O}_{V_0}\}^{F^{n+1}} \cdot \frac{dt^{F^{n+1}}}{t^{F^{n+1}}}$$

as desired. This completes the verification of the Lemma in this case.

Thus, it remains to deal with the nodes of X . Write $S = \text{Spec}(A)$. Let x be a node of X . To see that the isomorphism that we have already defined over the smooth locus extends, it suffices to work over the completion of X at x . Let $R = A[[s, t]]/(st)$; by abuse of notation, we shall denote the images of s and t in R by s and t , respectively. Since the deformation X^{\log} of X_0^{\log} is assumed to be equisingular, there

exists an isomorphism of the completion of X at x with $\mathrm{Spf}(R)$. Let us fix such an isomorphism. Note that Lemma 2.3 tells us that the radii of the monodromy operators of (P, ∇_P) at the nodes are elements of $(\mathbf{Z}/p^{n+1}\mathbf{Z})^\times$. Thus, by applying our analysis of the marked point case to each of the branches (and using the fact that sections of $\mathrm{Ad}(P)$ over $\mathrm{Spf}(R)$ are the same as pairs of sections of $\mathrm{Ad}(P)$ over each of the two branches whose values coincide at the node x), it follows that $\mathrm{Ad}(P)|_{\mathrm{Spf}(R)}$ is isomorphic to

$$(s^c + t^{-c})R \oplus R \oplus (s^{-c} + t^c)R$$

Thus, it suffices to compute what happens on the factor $(s^c + t^{-c})R$. But one computes easily that the cokernel in question is equal to

$$\{(s \cdot R_0)^{F^{n+1}} \oplus (k \cdot 1 + t \cdot R_0)^{F^{n+1}}\} \frac{dt^{F^{n+1}}}{t^{F^{n+1}}}$$

as desired. This completes the proof of the Lemma. \circ

From now on, we switch to a global point of view (on X). Thus, for the rest of this §, we assume that $X^{\log} \rightarrow S^{\log}$ is an r -pointed stable log-curve of genus g , and that (P, ∇_P) is torally indigenous on X^{\log} .

Corollary 2.5. *Under the hypotheses of Lemma 2.4, we have*

$$H^1(X_0, C^{\mathrm{nv}, [n+1]}) = H^1(X_0, C^{\mathrm{tr}, [n+1]}) = 0$$

and

$$H^0(X_0, K^{\mathrm{nv}, [n+1]}) = H^0(X_0, K^{\mathrm{tr}, [n+1]}) = 0$$

Moreover, the dimension over k of $H^0(X_0, C^{\mathrm{tr}, [n+1]})$ and $H^1(X_0, K^{\mathrm{tr}, [n+1]})$ is $3g - 3 + r$.

Proof. First, let us note that $H^0(X_0, K^{\mathrm{nv}, [n+1]}) = H^0(X_0, K^{\mathrm{tr}, [n+1]}) = 0$ follows from the fact that $\mathrm{Ad}(P_0)$ has no global sections. Next, let us show that $H^1(X_0, C^{\mathrm{nv}, [n+1]}) = H^1(X_0, C^{\mathrm{tr}, [n+1]}) = 0$. To do this, note that by Grothendieck duality on X_0 , it suffices to show that $\mathrm{Hom}_{X_0}(\mathcal{F}, \omega_{X_0^{F^{n+1}}}) = 0$, where \mathcal{F} is $C^{\mathrm{nv}, [n+1]}$ or $C^{\mathrm{tr}, [n+1]}$. Now observe that if $x \in X_0(k)$ is a node, \mathcal{O} is the local ring of X_0 at x , and \mathfrak{m} is its maximal ideal, then $\mathrm{Hom}_{\mathcal{O}}(\mathfrak{m}, \mathcal{O}) = \mathrm{Hom}_{\mathcal{O}}(\mathfrak{m}, \mathfrak{m})$ (i.e., sections of $\mathrm{Hom}_{\mathcal{O}}(\mathfrak{m}, \mathcal{O})$ are automatically zero at the node x). Thus, it follows from the explicit calculations in the proof of Lemma 2.4 that when one restricts a section of $\mathrm{Hom}_{X_0}(\mathcal{F}, \omega_{X_0^{F^{n+1}}})$ to a connected component Z of the normalization of X_0 (and tensors with $\omega_{X_0^{F^{n+1}}}^{-1}$), one obtains a horizontal morphism

$\mathrm{Ad}(P_0)|_Z \rightarrow \mathcal{O}_Z$, which must vanish, since (P_0, ∇_{P_0}) is crys-stable. This completes the proof of all the vanishing results.

With the vanishing results just proven, in order to compute the dimensions of the remaining cohomology groups, it suffices to compute the various Euler characteristics. Let Y be the normalization of X_0 . Let ν be the number of nodes of X_0 . Let $\mathcal{F} = K^{\mathrm{tr}, [n+1]}$; let \mathcal{G} be the quotient of $\mathcal{F}|_Z$ by its torsion. Then it follows from the proof of Lemma 2.4 that we have an injection $\mathcal{F} \hookrightarrow \mathcal{G}$ whose quotient is a torsion sheaf of length ν . Thus, $\chi(\mathcal{F}) = \chi(\mathcal{G}) - \nu$. Let us compute $\chi(\mathcal{G})$. Note that $\chi(\mathcal{G})$ is the sum of the $\chi(\mathcal{G}|_Z)$, over all connected components Z of Y . Let Z be a connected component of Y . Then it follows from the proof of Lemma 2.4 that $\mathcal{G}|_Z$ has degree $-s_Z$, where s_Z is the number of (former) nodes and marked points on Z . Thus, since $\mathcal{G}|_Z$ is a vector bundle of rank three on Z , $\chi(\mathcal{G}|_Z) = -3(g_Z - 1) - s_Z$ (where g_Z is the genus of Z). Summing, we obtain $-\chi(\mathcal{F}) = 3g - 3 + r$, as desired.

Finally, by Lemma 2.4, let us note that we have an injection

$$C^{\mathrm{tr}, [n+1]} \hookrightarrow K^{\mathrm{tr}, [n+1]} \otimes \omega_{(X_0^{\log})^{F^{n+1}}/k}$$

whose kernel is a torsion sheaf of length r . Thus, one can compute $\chi(C^{\mathrm{tr}, [n+1]})$ in terms of $\chi(K^{\mathrm{tr}, [n+1]})$; the result is $\chi(C^{\mathrm{tr}, [n+1]}) = 3g - 3 + r$, as desired. \circ

Corollary 2.6. *We retain the assumptions of Corollary 2.5. In particular, we suppose that $S_0 = \mathrm{Spec}(k)$, where k is a perfect field, and that the p^{n+1} -curvature \mathcal{P}_{n+1} of (P, ∇_P) is identically zero. Then $H^1 \stackrel{\mathrm{def}}{=} H_{\mathrm{DR}}^1(X, \mathrm{Ad}(P))$ (cf. Chapter I, §1.4) has a natural filtration F^i (for $0 \leq i \leq n+1$) such that $F^0 = H^1$; $F^i/F^{i+1} \cong H^0(X_0, C^{\mathrm{tr}, [i+1]})$ if $0 \leq i \leq n$.*

Proof. We would like to consider the spectral sequence associated to the hypercohomology of the toral de Rham complex $d^{\mathrm{tr}} : \mathrm{Ad}(P) \rightarrow \mathrm{Ad}^q(P) \otimes_{\mathcal{O}_X} \omega_{X^{\log}/S^{\log}}$ equipped with a certain filtration, defined as follows: Let C^\bullet denote the double complex obtained by considering d^{tr} over some Čech covering of X by two open affines. Let us assume that the first superscripted dot is the degree with respect to $d^{\mathrm{tr}} : \mathrm{Ad}(P) \rightarrow \mathrm{Ad}^q(P) \otimes_{\mathcal{O}_X} \omega_{X^{\log}/S^{\log}}$ (hence is 0 or 1), and that the second superscripted dot is the degree with respect to the Čech complex (hence is 0 or 1). We define $F^i(C^\bullet)$ (the i^{th} step of the filtration), where $i \geq 0$, to be the subcomplex of C^\bullet generated by $p^i \cdot C^{\bullet, 0}$ and $p^{\max(0, i-n-1)} \cdot C^{\bullet, 1}$. Then by Corollary 2.5, the E_{n+2}^\bullet -term of the spectral sequence consists only of $H^1(X_0, K^{\mathrm{tr}, [i]})$'s and $H^0(X_0, C^{\mathrm{tr}, [i]})$'s, all of which contribute only to H^1 . Moreover, the sums of the lengths of all these remaining terms is equal to the length of H^1 , so all the differentials of E_j^\bullet must vanish for $j \geq n+2$. Thus, the spectral sequence converges, and the resulting filtration on H^1 satisfies the necessary conditions. \circ

§2.3. Versal Families at Infinity

Now suppose that $S = \operatorname{Spec}(A)$, where A is a local artinian ring with perfect residue field k , and $J \subseteq A$ is a $\mathbf{Z}/p^{n+1}\mathbf{Z}$ -flat ideal such that $A/J = W(k)/p^{n+1}W(k)$. (Here $W(k)$ denotes the ring of Witt vectors with coefficients in k .) Let $I \subseteq J$ be an ideal annihilated by J such that, as an A/J -module, I is free of rank one. Let $T = \operatorname{Spec}(A/J)$; $Y^{\log} \rightarrow T^{\log}$ be the restriction of $X^{\log} \rightarrow S^{\log}$ to T ; let us assume that Y^{\log} is *equisingular*. Let us denote by a subscript I the result of restriction various objects over A to $\operatorname{Spec}(A/I)$. Let us assume that we are given a crys-stable bundle (P, ∇_P) (of radii ρ) on X^{\log} which is torally indigenous modulo I . Let (Q, ∇_Q) be the restriction of (P, ∇_P) to Y^{\log} . Let us assume that the pre- n -connection defined by $(P, \nabla_P)_I$ is a dormant n -connection.

Now as (P, ∇_P) ranges over all deformations of $(P, \nabla_P)_I$, we obtain a torsor over $H_I \stackrel{\text{def}}{=} I \otimes_A H_{\text{DR}}^1(Y, \operatorname{Ad}(Q))$ (which is a free A/J -module of rank $2(3g - 3 + r)$). Thus, by considering the deformation in the p -curvature resulting from a deformation of (P, ∇_P) , we obtain a morphism

$$\kappa_0 : H_I \rightarrow H^0(U_0, I_0 \otimes_A (\operatorname{Ad}(Q_0) \otimes_{\mathcal{O}_{Y_0}} (\Phi_Y)^* \omega_{Y^{\log}/T_0^{\log}})^{\text{hor}})$$

where $U \subseteq X$ is the complement of the nodes and marked points.

Now suppose that the pre- n -connection on (P, ∇_P) in fact defines a dormant j_0 -connection (for some $j_0 \leq n - 1$). Then we would like to define a filtration G^i (for $i = 0, \dots, j_0 + 2$) on H_I , as follows. Let $G^0 = H_I$. Let $G^1 \subseteq H_I$ be the kernel of the above morphism. Thus, deformations of (P, ∇_P) given by adding elements of G^1 have p -curvature zero, hence define 1-connections on (P, ∇_P) . Hence we can consider the deformation in the p^2 -curvature induced by such deformations. This gives a morphism

$$\kappa_1 : G^1 \rightarrow H^0(U_0, I_0 \otimes_A (\operatorname{Ad}(Q_0) \otimes_{\mathcal{O}_{Y_0}} (\Phi_Y^2)^* \omega_{Y^{\log}/T_0^{\log}})^{\text{hor}})$$

Then we let G^2 be the kernel of this morphism. Continuing in this fashion, we obtain morphisms

$$\kappa_i : G^i \rightarrow H^0(U_0, I_0 \otimes_A (\operatorname{Ad}(Q_0) \otimes_{\mathcal{O}_{Y_0}} (\Phi_Y^{i+1})^* \omega_{Y^{\log}/T_0^{\log}})^{\text{hor}})$$

for $i = 0, \dots, j_0 + 1$; hence G^i , for $i = 0, \dots, j_0 + 2$. Now we have the following result:

Lemma 2.7. *The filtration G^i (for $i = 0, \dots, j_0 + 2$) coincides with the filtration F^i of Corollary 2.6. Moreover, the morphisms $-\kappa_i$ are obtained by composing $G^i = F^i \rightarrow F^i/F^{i+1}$ with the isomorphism $F^i/F^{i+1} \cong H^0(X_0, C^{\text{tr}, [i+1]})$ of Corollary 2.6, followed by the isomorphism $C^{\text{tr}, [i+1]} \cong (\operatorname{Ad}^q(Q_0) \otimes_{\mathcal{O}_{Y_0}} (\Phi_Y^{i+1})^* \omega_{Y^{\log}/T_0^{\log}})^{\text{hor}}$ (of Lemma 2.4) and restriction from X_0 to U_0 .*

Proof. This lemma is a special case of the general principle “the derivative of the p -curvature is minus the Cartier operator,” and follows immediately from *Jacobson’s formula* (cf. the proof of Theorem 2.13 of [Mzk1], Chapter II). \square

Theorem 2.8. *Let g, r be such that $2g - 2 + r \geq 1$; let $n \geq 0$. Suppose that we are given radii $\rho_1, \dots, \rho_r \in (\mathbf{Z}/p^{n+1}\mathbf{Z})^\times$. Then the following objects exist:*

- (1) *a $\mathbf{Z}/p^{n+1}\mathbf{Z}$ -flat fine noetherian log scheme S^{\log} such that S_0 is smooth of dimension $(n+1)(3g-3+r)$ over \mathbf{F}_p ;*
- (2) *an r -pointed stable log-curve $X_0^{\log} \rightarrow S_0^{\log}$ of genus g such that the classifying morphism $S_0 \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ is smooth;*
- (3) *a crystal (P, ∇_P) on $\text{Crys}(X_0^{\log}/S^{\log})$ of radii ρ whose restriction to $\text{Crys}(X_0^{\log}/S_0^{\log})$ is totally indigenously and is such that (P, ∇_P) forms a dormant n -connection on (P_0, ∇_{P_0}) .*

Moreover, this collection of objects can be chosen to be versal at infinity in the following sense: There exists an open neighborhood W of the set of totally degenerate curves in $(\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ with the following property: Any other dormant n -connection on a totally indigenously bundle on an r -pointed stable log-curve of genus g over some $\mathbf{Z}/p^{n+1}\mathbf{Z}$ -flat log scheme T^{\log} whose classifying morphism modulo p maps into W is (étale locally on T) obtained by pull-back via some morphism $T^{\log} \rightarrow S^{\log}$.

Finally, when $n = 0$, we have the following: Let $\overline{\mathcal{N}}_{g,r}^{\rho, \text{dor}} \subseteq \overline{\mathcal{N}}_{g,r}^{\rho}$ denote the closed substack of dormant totally indigenously bundles of radii ρ . Then $\overline{\mathcal{N}}_{g,r}^{\rho, \text{dor}}$ is smooth and irreducible over \mathbf{F}_p of dimension $3g - 3 + r$ (if it is not empty), and finite and flat over $(\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$. Moreover, the pull-back to $\overline{\mathcal{N}}_{g,r}^{\rho, \text{dor}}$ of the divisor at infinity of $(\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ is a divisor with normal crossings, and $\overline{\mathcal{N}}_{g,r}^{\rho, \text{dor}} \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ is étale over points corresponding to totally degenerate curves.

Remark on Deformation Theory and Log Schemes. Before beginning the proof, we make a brief “general nonsense” remark about deformation theory and log schemes. Frequently, in the following proof, as well as in the remainder of the present work, we shall encounter the following situation: We have a base field k , a smooth scheme (or stack) M over k , and a morphism $S \rightarrow M$ of finite type. Moreover, M is equipped with a log structure arising from a divisor with normal crossings, and S is equipped with the log structure obtained by pulling back the log structure on M . Thus, we get a morphism of log schemes $S^{\log} \rightarrow M^{\log}$. Then observe that:

If S^{\log} is log smooth over k (where $\text{Spec}(k)$ is equipped with the trivial log structure), then S is smooth over k . Similarly, if $S^{\log} \rightarrow M^{\log}$ is log étale, then $S \rightarrow M$ is étale. Moreover, if S^{\log} is log smooth over k , then the

rank of the tangent bundle on S is always equal to the rank of the logarithmic tangent bundle on S^{\log} .

These observations follow immediately from the standard definitions of “smooth,” “log smooth,” “étale,” and “log étale” involving the existence (and uniqueness) of liftings of points to nilpotent thickenings, plus the fact that M is k -smooth. Indeed, if T is the spectrum of an artinian ring; $T_0 \subseteq T$ is defined by a nilpotent ideal; and one is given a morphism $T_0 \rightarrow S$, then one gets a morphism $T_0 \rightarrow S \rightarrow M$ which lifts to a morphism $T \rightarrow M$ (since M is k -smooth). Moreover, by pulling back the log structure on M to T , one gets a log morphism $T_0^{\log} \rightarrow S^{\log}$, together with an exact closed immersion $T_0^{\log} \subseteq T^{\log}$. Thus, if S^{\log} is log smooth over k , we get a lifting $T^{\log} \rightarrow S^{\log}$, hence a morphism $T \rightarrow S$, as desired. This shows that the log smoothness of S^{\log} implies the smoothness of S . The assertion on étale morphisms follows by a similar (but even easier) argument. The assertion on the rank of the tangent bundles follows from the fact that if S^{\log} is log smooth over k , then it is necessarily smooth over k on an open dense subset of S . The reason that we make these observations is that often our computations of de Rham cohomology will make it easier to show directly (using the formalism of deformation theory) that S^{\log} is log smooth or that $S^{\log} \rightarrow M^{\log}$ is log étale, but in the following, we will immediately conclude from such arguments (without further comment) that S is smooth or that $S \rightarrow M$ is étale.

Proof. Let us first consider the case $n = 0$. In the above discussion, the p -curvature of (P, ∇_P) defines an element of $H^0(X_0, I_0 \otimes_A (\mathrm{Ad}^q(Q_0) \otimes_{\mathcal{O}_{Y_0}} (\Phi_Y)^* \omega_{Y_0^{\log}/T_0^{\log}})^{\mathrm{hor}})$. But, by Lemma 2.7, H_I surjects onto this H^0 . Thus, in other words, the obstruction to deforming a dormant torally indigenous bundle vanishes. Moreover, the dormant infinitesimal deformations of $(P_0, \nabla_{P_0})_I$ form a torsor over the kernel of $(H_I)_{\mathbf{F}_p} \rightarrow F^0/F^1$, which has k -dimension $3g - 3 + r$ (cf. Corollaries 2.5, 2.6). Thus, it follows that $\overline{\mathcal{N}}_{g,r}^{\rho, \mathrm{dor}}$ is smooth, of dimension $3g - 3 + r$ (if it is not empty). Since, by Theorem 1.12, it is also geometrically connected, it follows that it is geometrically irreducible. Since $\overline{\mathcal{N}}_{g,r}^{\rho, \mathrm{dor}}$ is finite over $\overline{\mathcal{M}}_{g,r}$ and regular of the same dimension as $\overline{\mathcal{M}}_{g,r}$, it follows from commutative algebra that it is flat over $\overline{\mathcal{M}}_{g,r}$. Finally, to see that $\overline{\mathcal{N}}_{g,r}^{\rho, \mathrm{dor}} \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ is étale over points corresponding to totally degenerate curves, it suffices to observe that the fiber of this morphism over such points is *reduced*. But this follows from Lemma 2.3: Indeed, Lemma 2.3 implies that the radii at the nodes (of a dormant torally indigenous bundle on a totally degenerate curve) lie in \mathbf{F}_p (which is a reduced ring!); moreover, once the radii at all the marked points and nodes of a torally indigenous bundle on a totally degenerate curve are determined, the torally indigenous bundle itself is completely determined (cf. Chapter I, Theorem 4.4); thus, the fiber in question is reduced. This completes the proof in the case $n = 0$.

Now we assume that $n \geq 1$, and that the case $n - 1$ has been

completed, and we consider the case n . Let S' and $X'_0 \rightarrow S'_0$ be the “ S ” and “ $X_0 \rightarrow S_0$ ” obtained in the case $n - 1$. Then by taking a torsor $T_0 \rightarrow S'_0$ over the first de Rham cohomology module of $\mathrm{Ad}(P'_0)$, and lifting T_0 arbitrarily to some $\mathbf{Z}/p^{n+1}\mathbf{Z}$ -flat T , one sees that one can construct a versal family (at infinity) $(h^{\log} : Y_0^{\log} \rightarrow T_0^{\log}; (Q, \nabla_Q))$ of torally indigenous bundles equipped with n -connections. (Here, we may take Y_0^{\log} , the log structure of T^{\log} , and $(Q_{n-1}, \nabla_{Q_{n-1}})$ to be the pull-backs of the corresponding objects over S' .) We wish to consider the p^{n+1} -curvature of (Q, ∇_Q) . Let $V_0 \subseteq Y_0$ be the complement of the marked points and nodes. Let us consider the vector bundle $\mathcal{G} \stackrel{\mathrm{def}}{=} \mathrm{Ad}(Q_0) \otimes_{\mathcal{O}_{Y_0}} (\Phi_Y^{n+1})^* \omega_{Y_0^{\log}/T_0^{\log}}$ on Y_0 . Note that since $h_* \omega_{Y_0^{\log}/T_0^{\log}}^{\otimes(j-p^{n+1})} = 0$ (for $j = 0, 1, 2$), we have $\mathbf{R}^1 h_* \mathcal{G} = 0$. Thus, $h_* \mathcal{G}$ is a vector bundle on T_0 . The p^{n+1} -curvature \mathcal{P}_{n+1}^V of (Q, ∇_Q) over V_0 is (by definition) a section of \mathcal{G} over V_0 . On the other hand, it follows from the explicit calculation in the proof of Lemma 2.3 that over a generic point η of T_0 , \mathcal{P}_{n+1}^V extends over the marked points. But the union of V_0 and the marked points forms an open subset of Y_0 whose complement is of codimension ≥ 2 . Thus, we obtain a section \mathcal{P}_{n+1} of \mathcal{G} over Y_0 . By abuse of notation, we shall regard \mathcal{P}_{n+1} as a section of the vector bundle $h_* \mathcal{G}$.

Let $S_0 \subseteq T_0$ be the zero locus of the section \mathcal{P}_{n+1} of $h_* \mathcal{G}$. Then it follows from Lemma 2.7 that S_0 is smooth of dimension $3g - 3 + r$ over S'_0 . After possibly replacing T by some T' where $T' \rightarrow T$ is étale surjective, we may assume that $S_0 \subseteq T_0$ lifts to a closed $\mathbf{Z}/p^{n+1}\mathbf{Z}$ -flat subscheme $S \subseteq T$. Then by restricting T^{\log} , $Y_0^{\log} \rightarrow T_0^{\log}$, and (Q, ∇_Q) to S , we obtain S^{\log} , $X_0^{\log} \rightarrow S_0^{\log}$, and (P, ∇_P) . These data satisfy all the conditions in the statement of the Theorem, except possibly (2). That is, it remains to see that the classifying morphism $S_0 \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ is smooth. In fact, it may not be smooth, but we know that at least the fiber of S_0 over a point of $(\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ corresponding to a totally degenerate curve is smooth over \mathbf{F}_p . (Indeed, this follows by induction on n : We know the result in the case $n = 0$ (by what was done in the first paragraph of this proof); the “induction step” then follows from the fact that S_0 is smooth over S'_0 .) Thus, by replacing S by some open subscheme of S , we obtain versal objects at infinity, as stated in the Theorem. This completes the induction step, and hence the proof of the Theorem. \circ

Corollary 2.9. *Let g, r be such that $2g - 2 + r \geq 1$; let $n \geq 1$. Suppose that we are given radii $\rho_1, \dots, \rho_r \in (\mathbf{Z}/p^{n+1}\mathbf{Z})^\times$. Then the following objects exist:*

- (1) *a $\mathbf{Z}/p^{n+1}\mathbf{Z}$ -flat fine noetherian log scheme S^{\log} such that S_0 is smooth of dimension $(n + 2)(3g - 3 + r)$ over \mathbf{F}_p ;*
- (2) *an r -pointed stable log-curve $X_0^{\log} \rightarrow S_0^{\log}$ of genus g such that the classifying morphism $S_0 \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ is smooth;*
- (3) *a crystal (P, ∇_P) on $\mathrm{Crys}(X_0^{\log}/S^{\log})$ of radii ρ whose restriction to $\mathrm{Crys}(X_0^{\log}/S_0^{\log})$ is torally indigenous and is such that (P, ∇_P) forms an n -connection on (P_0, ∇_{P_0}) .*

Moreover, this collection of objects can be chosen to be versal at infinity in the following sense: There exists an open neighborhood W of the set of totally degenerate curves in $(\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ with the following property: Any other n -connection on a torally indigenous bundle on an r -pointed stable log-curve of genus g over some $\mathbf{Z}/p^{n+1}\mathbf{Z}$ -flat log scheme T^{\log} whose classifying morphism modulo p maps into W is (étale locally on T) obtained by pull-back via some morphism $T^{\log} \rightarrow S^{\log}$.

Proof. One takes for the S in the present Corollary the “ T ” of the proof of Theorem 2.8. \square

Let S^{\log} be a $\mathbf{Z}/p^{n+1}\mathbf{Z}$ -flat fine noetherian log scheme; $f_0^{\log} : X_0^{\log} \rightarrow S_0^{\log}$ be an r -pointed stable log-curve of genus g ; and (P, ∇_P) an n -connection of radii ρ on a torally indigenous bundle on X_0^{\log} . Let us assume that X_0^{\log} is “close to infinity” in the sense that its classifying morphism maps into the W of Corollary 2.9. Let $U_0 \subseteq X_0$ be the complement of the marked points. Let $\mathcal{F} \stackrel{\text{def}}{=} \text{Ad}(P_0) \otimes_{\mathcal{O}_{X_0}} (\Phi_X^{n+1})^* \omega_{X_0^{\log}/S_0^{\log}}$ on X_0 . As we saw in the proof of Theorem 2.8, the p^{n+1} -curvature \mathcal{P}_{n+1} of (P, ∇_P) is a section of \mathcal{F} over X_0 whose restriction to V_0 is annihilated by the successive connections $\nabla_{\mathcal{A}^{[0]}}, \dots, \nabla_{\mathcal{A}^{[n]}}$. (In fact, we only saw this for the “ S ” in Theorem 2.8, but that “ S ” was versal, so the claim that \mathcal{P}_{n+1} is defined over all of X_0 is valid for all “ S ” that are “close to infinity.”) By using the Killing form $\langle -, - \rangle$ on $\text{Ad}(P_0)$, we can form the section $\delta \stackrel{\text{def}}{=} \langle \mathcal{P}_{n+1}, \mathcal{P}_{n+1} \rangle$ of $(\Phi_X^{n+1})^* \omega_{X_0^{\log}/S_0^{\log}}^{\otimes 2}$ over X_0 . It is not difficult to see that δ vanishes at the marked points. Moreover, since the Killing form is horizontal for all the connections involved, it follows that δ in fact descends to a section ζ over X_0 of

$$(\omega_{X_0^{\log}/S_0^{\log}}^{\otimes 2}(-M_f))^{F^{n+1}}$$

(where $M_f \subseteq X_0$ is the divisor of marked points). Let $Q_S^{n+1} \rightarrow S$ be the geometric vector bundle associated to the locally free sheaf

$$f_*(\omega_{X_0^{\log}/S_0^{\log}}^{\otimes 2}(-M_f))^{F^{n+1}}$$

on S_0 . Then we have an $(n+1)$ -Verschiebung morphism over S_0 :

$$\mathcal{V}_S^{n+1} : S_0 \rightarrow Q_S^{n+1}$$

defined by ζ . Let

$$\mathcal{N}_S \subseteq S_0$$

be the closed subscheme defined by considering the zero locus of ζ . Thus, \mathcal{N}_S is defined locally by $3g - 3 + r$ equations. Moreover, over \mathcal{N}_S , the connection $\nabla_{\mathcal{A}^{[n]}}$ on the bundle $\mathcal{A}^{[n]}$ will have nilpotent p -curvature.

Finally, we propose to show that by (essentially) the same argument as that given for Proposition 1.8, one can prove the following analogue (Proposition 2.10) of Proposition 1.8: Suppose (in the above discussion) that $X_0^{\log} \rightarrow S^{\log}$ and (P, ∇_P) are versal at infinity, as in Corollary 2.9. Moreover, let T^{\log} be a $\mathbf{Z}/p^{n+1}\mathbf{Z}$ -flat lifting of the “ S^{\log} ” for $n-1$ of Theorem 2.8. Thus, we have a morphism $S^{\log} \rightarrow T^{\log}$ such that $S \rightarrow T$ is smooth of dimension $2(3g-3+r)$; $X_0^{\log} \rightarrow S_0^{\log}$ is pulled back from a curve $Y_0^{\log} \rightarrow T_0^{\log}$; and $(P, \nabla_P)_{n-1}$ is pulled back from $\text{Crys}(Y_0^{\log}/T_{n-1}^{\log})$.

Let $A \stackrel{\text{def}}{=} W(k)/p^{n+1} \cdot W(k)$ (where k is perfect of characteristic p). Let $\xi \in T(A_{n-1})$ be a point such that the restriction of $(P, \nabla_P)_{n-1}$ to ξ defines a torally indigenous bundle on a *totally degenerate curve* over $\text{Spec}(A_{n-1})^{\log}$ (where we endow $\text{Spec}(A_{n-1})$ with the log structure pulled back via ξ from T). Let $\xi^\rho \in \mathcal{N}_S(k)$ be a point lying over ξ_0 . Now we want to consider the morphism $\mathcal{N}_S \subseteq S_0 \rightarrow T_0$.

Proposition 2.10. *In the above notation, there exists a complete local ring R_ξ which is formally smooth over the complete local ring $\widehat{\mathcal{O}}_{T_0, \xi_0}$ of relative dimension $3g-3+r$ and which is equipped with an $\widehat{\mathcal{O}}_{T_0, \xi_0}$ -algebra morphism $\psi_\xi : R_\xi \rightarrow \widehat{\mathcal{O}}_{\mathcal{N}_S, \xi^\rho}$ such that ψ_ξ is a finite, flat local complete intersection morphism of degree $\leq 4^{3g-3+r}$. In particular, if $p > 4^{3g-3+r}$, then every irreducible component of $(\mathcal{N}_S)_{\text{red}}$ that passes through ξ^ρ is generically smooth of relative dimension $3g-3+r$ over T_0 .*

Proof. Just as Proposition 1.8 followed essentially from the computation of the determinant of the p -curvature in §1.2, the key to proving this result is the computation of the determinant of the p^{n+1} -curvature. Thus, let B^{\log} be a log scheme such that $B = \text{Spec}(A)$. Let $Z^{\log} \rightarrow B^{\log}$ be a totally degenerate r -pointed stable curve of genus g . Let (Q, ∇_Q) be a torally indigenous bundle on Z^{\log} (of radii ρ) that defines an n -connection on (Q_0, ∇_{Q_0}) . Note that Z^{\log} has precisely $3g-3+r$ nodes. The determinant of the p^{n+1} -curvature of (Q, ∇_Q) is thus determined by its values on the nodes of Z_0 . Thus, it suffices to compute the determinant of the p^{n+1} -curvature of (Q, ∇_Q) at a node of Z_0 .

Let $W^{\log} \rightarrow B^{\log}$ be one of the 3-pointed curves of genus 0 over B^{\log} that are used to construct Z^{\log} (by clutching, as in the discussion of Chapter I, §3.3). Let $W^{\log} \hookrightarrow Z^{\log}$ be the natural inclusion morphism. We would like to see how the determinant of the p^{n+1} -curvature of the restriction of (Q, ∇_Q) to W^{\log} at a marked point ν of W that maps to a node of Z varies as the connection ∇_Q varies. Here, when we say that ∇_Q “varies,” in fact, we mean that it is fixed modulo p^n ; that is to say, we only wish to vary the lifting of ∇_Q to a connection over $\mathbf{Z}/p^{n+1}\mathbf{Z}$.

Let us write $(Q, \nabla_Q)|_{W^{\log}}$ in a neighborhood of ν as the projectivization of a rank two vector bundle with connection $(\mathcal{E}, \nabla_{\mathcal{E}})$ whose determinant is trivial. Write $(\mathcal{F}, \nabla_{\mathcal{F}})$ for the vector bundle $t^a \cdot \mathcal{E}$ (equipped with the connection $\nabla_{\mathcal{F}}$ naturally induced by $\nabla_{\mathcal{E}}$), where a is as in

the second to last paragraph of the proof of Lemma 2.3. Choose local sections s_1 and s_2 of \mathcal{F} (in a neighborhood of ν) such that (s_1, s_2) defines a horizontal isomorphism of $\mathcal{O}_{W_{n-1}} \oplus (t^{a-b} \cdot \mathcal{O}_{W_{n-1}})$ with $(\mathcal{F}, \nabla_{\mathcal{F}})_{n-1}$ (where $b > a$ is as in Lemma 2.3). Let us write $\mathcal{G}_{n-1} \subseteq \mathcal{F}_{n-1}$ for the \mathcal{O}_W -submodule of \mathcal{F}_{n-1} generated by s_1 and s_2 . For $i = 1, 2$, let σ_i be the result of dividing $\nabla_{\mathcal{F}}(s_i)$ by p^n and reducing modulo p . Thus, σ_i is a section of $\mathcal{F}_0 \otimes_{\mathcal{O}_{W_0}} \omega_{W_0^{\log}}$. If we then apply the Cartier operator a total of n times to σ_i (for $i = 1, 2$), we get a section δ_i of

$$\mathcal{G}_0 \otimes_{\mathcal{O}_{W_0}} (\Phi_W^n)^* \omega_{(W_0^{\log})^{F^n}} \subseteq \mathcal{F}_0 \otimes_{\mathcal{O}_{W_0}} (\Phi_W^n)^* \omega_{(W_0^{\log})^{F^n}}$$

If we then write the values of δ_i (for $i = 1, 2$) at ν relative to the basis of \mathcal{G}_0 defined by s_1 and s_2 , we obtain a two-by-two matrix μ'_ν with entries in k . Let $\mu_\nu \stackrel{\text{def}}{=} \mu'_\nu - \frac{1}{2} \text{tr}(\mu'_\nu)$ (where “tr” denotes the trace of the matrix). Sorting through the definitions, we see that

$$\begin{aligned} \det(\mu_\nu^p - \mu_\nu) &= (\det(\mu_\nu)^{\frac{p-1}{2}} - 1)^2 \det(\mu_\nu) \\ &= \det(\mu_\nu) \cdot \prod_{\lambda \in (\mathbf{F}_p^\times)^2} (\det(\mu_\nu) - \lambda)^2 \end{aligned}$$

is the value of the determinant of the p^{n+1} -curvature of (Q, ∇_Q) at ν .

Now as the connection ∇_Q varies (while being held fixed modulo p^n), it is clear that μ'_ν varies *linearly* in the affine moduli of the connection ∇_Q . Thus, μ_ν also varies linearly in the affine moduli of ∇_Q , and hence $\det(\mu_\nu)$ varies quadratically in the affine moduli of ∇_Q . Thus, in summary, we conclude that (in its dependence on ∇_Q) *the value of the determinant of the p^{n+1} -curvature of (Q, ∇_Q) at ν is a product of relatively prime polynomials of degree ≤ 4* . Moreover, since by Lemma 2.3, (1), (and Chapter I, Theorem 4.4) there are only finitely many nilpotent n -connections on (Q_0, ∇_{Q_0}) that give rise to torally indigenous bundles on Z^{\log} , it is then a matter of elementary commutative algebra (and “Bézout’s Theorem” in the form of Lemma 2.11 below, where we take “ n ” to be $3g - 3 + r$ and “ d ” to be 4) to draw the conclusions stated in Proposition 2.10. Note that the upper estimate of 4^{3g-3+r} arises from the upper estimate of 4 (for the degree of the prime factors of the value of the determinant of the p^{n+1} -curvature at ν) just derived above, combined with the fact that there are a total of $3g - 3 + r$ nodes. Finally, observe that the extra $3g - 3 + r$ dimensions of R_ξ over $\widehat{\mathcal{O}}_{T_0, \xi_0}$ arise from the fact that in the above discussion, we only varied the connection ∇_Q , assuming all the time that (Q, ∇_Q) is *torally indigenous* on Z^{\log} . \circ

Lemma 2.11. *Let k be an algebraically closed field. Let n and d be positive integers. Let $R \stackrel{\text{def}}{=} k[x_1, \dots, x_n]$; $U \stackrel{\text{def}}{=} \text{Spec}(R)$. Let $f_1, \dots, f_n \in R$ be n polynomials of degree $\leq d$. Let $F \stackrel{\text{def}}{=} V(f_1, \dots, f_n) \subseteq U$. Suppose that the origin $0_n \stackrel{\text{def}}{=} (0, \dots, 0) \in U(k)$ is a connected component of F_{red} . Then the length of the connected component of F with support at 0_n is $\leq d^n$.*

Proof. Let P be projective n -space over k . Thus, U embeds naturally as an open subscheme of P , and we may think of f_1, \dots, f_n as defining sections s_1, \dots, s_n of $\mathcal{O}_P(d)$. The lemma then follows from “Bézout’s Theorem,” as given in [Fulton1], p. 226, Example 12.3.7, (iii). \square

Remark. The reader who is disturbed by the lack of naturality of the schemes “ S ” constructed in Theorem 2.8 and Corollary 2.9 – i.e., the fact that they are, for instance, only *versal at infinity*, and not *universal*, and do not represent any natural functor – will thus appreciate the approach of Chapter III to parametrizing this sort of data. This approach gives rise to entirely natural algebraic stacks, but one pays a price in that the approach of Chapter III, i.e., VF-patterns, does not allow one see what happens when the curve is singular or has (toral) marked points. Indeed, we shall see in Chapter VI that relative to the point of view of VF-patterns, there is no place for dormant indigenous bundles on singular or marked curves. On the other hand, the author sees no obvious way to construct or show the nonemptiness of stacks of VF-patterns with zeroes except by using degeneration to singular curves and then applying the techniques of the present §. Perhaps the reader will see a simple way to unify these two disparate points of view.

§3. Mildly Spiked Bundles

Whereas in §2, we studied (in particular) bundles whose p -curvature vanishes identically, in this §, we study bundles whose p -curvature is nilpotent and has zeroes, but is not identically zero. In particular, we will be concerned with the deformation theory of such bundles on singular curves, since it is this deformation theory that will allow us in Chapter IV, §3, to construct explicit examples of such bundles on smooth curves.

§3.1. Definition and First Properties

Let $f^{\log} : X^{\log} \rightarrow S^{\log}$ be a locally stable morphism of dimension one, where S^{\log} is a fine noetherian log scheme of characteristic $p > 2$. Let (P, ∇_P) be a \mathbf{P}^1 -bundle with connection on X^{\log} . Let us think of its p -curvature of (P, ∇_P) as a horizontal morphism

$$\mathcal{P} : \Phi_X^* \tau_{X^{\log}/S^{\log}} \rightarrow \mathrm{Ad}(P)$$

Let us assume that the image of \mathcal{P} is *square nilpotent*, and that (P, ∇_P) has *nonzero monodromy* at the nodes of fibers of $f : X \rightarrow S$. Let $V_{\mathcal{P}} \subseteq X$ be the zero locus of \mathcal{P}^{\vee} . Thus, $V_{\mathcal{P}} \subseteq X$ is a closed subscheme whose defining ideal $\mathcal{I}_{\mathcal{P}}$ is stabilized by the logarithmic exterior derivative on \mathcal{O}_X .

Definition 3.1. We shall call $V_{\mathcal{P}}$ the *spiked locus* of (P, ∇_P) . We shall call (P, ∇_P) *mildly spiked (of strength d)* if $V_{\mathcal{P}}$ is an S -flat relative divisor (of degree d).

Proposition 3.2. Suppose that $X \rightarrow S$ is proper. Then there exists a subscheme $Z \subseteq S$ such that for any morphism $T^{\log} \rightarrow S^{\log}$, the pull-back to $X^{\log} \times_{S^{\log}} T^{\log}$ of (P, ∇_P) is mildly spiked of strength d if and only if $T \rightarrow S$ factors through Z .

Proof. By means of the theory of flattening stratifications, we may assume that $V_{\mathcal{P}}$ is finite and flat over S of degree d . Moreover, being a divisor is an open condition on X . Since $X \rightarrow S$ is proper, this completes the proof. \circ

For instance, the locus of mildly spiked torally indigenous bundles of strength d and radii ρ on an r -pointed stable curve of genus g (where $2g - 2 + r \geq 1$) forms an algebraic substack

$$\overline{\mathcal{N}}_{g,r}^{\rho}[d] \subseteq \overline{\mathcal{N}}_{g,r}^{\rho}$$

The goal of this § is to show that this substack is smooth of dimension $3g - 3 + r$ (everywhere), and étale over $\overline{\mathcal{M}}_{g,r}$ at the points where X^{\log} is totally degenerate.

Suppose that S is the spectrum of a field k . Let $\nu \in X(k)$ be a node. Let X_{ν} be the completion of X at ν ; let Y_{ν} and Z_{ν} be the two irreducible components of X_{ν} . Suppose that the $V_{\mathcal{P}}$ is of finite length $l > 0$ at ν . Let (Q, ∇_Q) (respectively, (R, ∇_R)) be the restriction of (P, ∇_P) to Y_{ν} (respectively, Z_{ν}). Let \mathcal{P}_Q (respectively, \mathcal{P}_R) be the p -curvature of (Q, ∇_Q) (respectively, (R, ∇_R)). Then \mathcal{P}_Q (respectively, \mathcal{P}_R) factors through a line bundle $\mathcal{L} \subseteq \text{Ad}(Q)$ (respectively, $\mathcal{M} \subseteq \text{Ad}(R)$) such that the inclusions $\mathcal{L} \hookrightarrow \text{Ad}(Q)$, $\mathcal{M} \hookrightarrow \text{Ad}(R)$ are locally split. Note that since (P, ∇_P) is nilpotent and S is reduced, the radius of (Q, ∇_Q) and (R, ∇_R) at ν is some $\rho \in \mathbf{F}_p^{\times}$.

Let y (respectively, z) be a local uniformizer on Y_{ν} (respectively, Z_{ν}) that vanishes at ν . Let $\theta_y = y\partial/\partial y$, $\theta_z = z\partial/\partial z$. Moreover, the monodromy operator μ evaluated on θ_y (respectively, θ_z) acts on $\text{Ad}(Q)|_{\nu}$ (respectively, $\text{Ad}(R)|_{\nu}$) with eigenvalues 0 and $\pm 2\rho$. Since the subspace $\mathcal{L}|_{\nu}$ of $\text{Ad}(Q)|_{\nu}$ is clearly nilpotent and stabilized by the monodromy operator (this follows since \mathcal{P} is horizontal and nilpotent), we thus conclude that the subspace $\mathcal{L}|_{\nu}$ is the eigenspace of $\mu(\theta_y)$ of eigenvalue a , where $a = \pm 2\rho$. Similarly, $\mathcal{M}|_{\nu}$ is the eigenspace of $\mu(\theta_z)$ of eigenvalue b , where $b = \pm 2\rho$. Then since the monodromy operator on $\Phi_X^* \tau_{X^{\log}/S^{\log}}$ is zero, it follows that the length l_Q (respectively, l_R) of the zero locus of \mathcal{P}_Q (respectively, \mathcal{P}_R) is $\equiv -a \pmod{p}$ (respectively, $\equiv -b \pmod{p}$).

Suppose that $a = b$. Note that $\mu(\theta_y) = -\mu(\theta_z)$ (as operators on $\text{Ad}(P)|_{\nu}$). Thus, it follows that \mathcal{P}^{\vee} factors as a surjection $\text{Ad}(P) \rightarrow \mathcal{L}^{\vee} \oplus \mathcal{M}^{\vee}$ followed by an injection $\mathcal{L}^{\vee} \oplus \mathcal{M}^{\vee} \hookrightarrow \Phi_X^* \omega_{X^{\log}/S^{\log}}$. In particular, $V_{\mathcal{P}}$ cannot be a divisor at ν .

Now, suppose that $a \neq b$, i.e., $l_Q + l_R \equiv 0 \pmod{p}$. Then it follows that \mathcal{L} and \mathcal{M} glue together to form a line bundle \mathcal{H} on X , equipped with a locally split injection $\mathcal{H} \hookrightarrow \mathrm{Ad}(P)$. Moreover, \mathcal{P}^\vee factors through the dual of this injection $\mathrm{Ad}(P) \rightarrow \mathcal{H}^\vee$. That is to say, $V_{\mathcal{P}}$ is defined at ν as the zero locus of a morphism of line bundles $\mathcal{H}^\vee \rightarrow \Phi_X^* \omega_{X^{\log}/S^{\log}}$ on X_ν . In particular, it follows that (P, ∇_P) will be mildly spiked at ν . Moreover, the degree of $V_{\mathcal{P}}$ at ν will be divisible by p . In summary, we have the following

Proposition 3.3. *Let S^{\log} be a fine noetherian scheme of characteristic $p > 2$. Let (P, ∇_P) be a mildly spiked bundle on X^{\log} , and suppose that ν is a node of a geometric fiber of $X \rightarrow S$. Then the length of $V_{\mathcal{P}}$ at ν is divisible by p .*

§3.2. De Rham Cohomology Computations

Let $f^{\log} : X^{\log} \rightarrow S^{\log}$, and (P, ∇_P) be as in the preceding subsection. Moreover, let us assume that S is the spectrum of an algebraically closed field k , and that we are given a horizontal surjection

$$\mathrm{Ad}(P) \rightarrow \mathcal{W}$$

(where \mathcal{W} is a line bundle) such that the image of its dual $\mathcal{W}^\vee \subseteq \mathrm{Ad}(P)$ is square nilpotent. Thus, \mathcal{W} is equipped with a logarithmic connection. Moreover, the p -curvature of $(\mathcal{W}, \nabla_{\mathcal{W}})$ is necessarily nilpotent, hence zero.

Let us call the points of X where $(\mathcal{W}, \nabla_{\mathcal{W}})$ has nonzero monodromy \mathcal{W} -active. Thus, \mathcal{W} -active points are either nodes or marked points. Let Z be the curve obtained from X by normalizing X at the \mathcal{W} -active nodes. We take for the marked points of Z the points that map to marked points or \mathcal{W} -active nodes of X . We denote the divisor of marked points on Z by $M_Z \subseteq Z$. Thus, we obtain a natural log structure on Z ; call the resulting log scheme Z^{\log} . Let $M_Z^{\mathrm{ac}} \subseteq Z$ denote the divisor of marked points of Z that map to \mathcal{W} -active points of X .

A typical (though by no means the only useful) example of a \mathcal{W} as above is the following: If (P, ∇_P) is mildly spiked of strength d , then $V_{\mathcal{P}} \subseteq X$ is a relative divisor over S of degree d . Thus, the dual \mathcal{P}^\vee of the p -curvature gives a horizontal surjection $\mathrm{Ad}(P) \rightarrow \mathcal{W}$, where $\mathcal{W} = (\Phi_X^* \omega_{X^{\log}/S^{\log}})(-V_{\mathcal{P}})$. We shall refer to this case by saying that \mathcal{W} is mild. In the mild case, let us denote by $V_Z \subseteq Z$ the schematic inverse image of V_X in Z (via the normalization morphism $Z \rightarrow X$). Let us denote by $D_Z \subseteq Z^F$ the subscheme (divisor) defined by the ideal which is the kernel of the morphism

$$\mathcal{O}_{Z^F} \rightarrow \mathcal{O}_Z \rightarrow \mathcal{O}_{V_Z}$$

where the first map is the Frobenius on Z , and the second map is the natural surjection.

Returning to the general (i.e., not necessarily mild) case, let us denote by $M_X \subseteq X$ the divisor of marked points on X ; by $M_X^{\text{non}} \subseteq X$ the divisor of marked points that are *not* \mathcal{W} -active; and by $M_X^{\text{ac}} \subseteq X$ the divisor of \mathcal{W} -active marked points. Thus, $\omega_{X^{\log}/S^{\log}}(-M_X^{\text{non}})|_Z = \omega_{Z/k}(M_Z^{\text{ac}})$. Now we have a complex:

$$\mathcal{W} \xrightarrow{\nabla_{\mathcal{W}}} \mathcal{W} \otimes_{\mathcal{O}_X} \omega_{X^{\log}/S^{\log}}(-M_X^{\text{non}})$$

whose first hypercohomology module we denote by $H_{\text{DR}}^1(X, \mathcal{W})$. Let us denote the kernel (respectively, cokernel) of this complex by \mathcal{K} (respectively, \mathcal{C}). Note that \mathcal{K} and \mathcal{C} have the natural structure of coherent sheaves over \mathcal{O}_{X^F} . In fact, \mathcal{K} even has the structure of a line bundle on Z^F .

We would like to study $H_{\text{DR}}^1(X, \mathcal{W})$. First note that by considering the spectral sequence whose E_2^* -term consists of the modules $H^i(X^F, \mathcal{C})$ and $H^i(X^F, \mathcal{K})$, we obtain an exact sequence

$$0 \rightarrow H^1(X^F, \mathcal{K}) \rightarrow H_{\text{DR}}^1(X, \mathcal{W}) \rightarrow H^0(X^F, \mathcal{C}) \rightarrow 0$$

Now we would like to consider \mathcal{K} and \mathcal{C} in more detail.

Proposition 3.4. (i) We have $\mathcal{C} \cong \mathcal{K} \otimes_{\mathcal{O}_{Z^F}} (\omega_{Z/k}(M_Z^{\text{ac}}))^F$. If \mathcal{W} is mild, then, $\mathcal{K} \cong (\omega_{Z^F/k}(-D_Z + M_Z^F))$.

(ii) Suppose that $X^{\log} \rightarrow S^{\log}$ is an r -pointed stable curve of genus g ; (P, ∇_P) is crys-stable; and \mathcal{W} is mild. Then $H^1(X^F, \mathcal{C}) = 0$; $H^0(X^F, \mathcal{C})$ has dimension over k equal to

$$3g - 3 + r - \deg(D_Z) + \deg(M_Z^{\text{ac}}) - N_X^{\text{ac}}$$

(where N_X^{ac} is the number of \mathcal{W} -active nodes on X); and $H^1(X^F, \mathcal{K})^\vee = \Gamma(Z^F, \mathcal{O}_{Z^F}(D_Z - M_Z^F))$.

Proof. (i) Direct computation using the Cartier isomorphism gives

$$\mathcal{C} \cong \mathcal{K} \otimes_{\mathcal{O}_{X^F}} (\omega_{X^{\log}/S^{\log}}(-M_X^{\text{non}}))^F$$

In the mild case, the fact that \mathcal{K} is as described also follows from direct computation.

(ii) Note that $\omega_{Z^F/k}(-D_Z + M_Z^F + (M_Z^{\text{ac}})^F)$ has degree > 0 on every irreducible component of Z^F . Indeed, this follows from the fact that $\mathcal{P}^\vee|_Z$ factors through $\Phi_Z^*(\omega_{Z^F/k}(-D_Z + M_Z^F + (M_Z^{\text{ac}})^F))$, and (P, ∇_P) is crys-stable. It thus follows immediately that every morphism $\mathcal{C} \rightarrow \omega_{Z^F/k}$ must vanish, hence that $H^1(Z^F, \mathcal{C}) = H^1(X^F, \mathcal{C}) = 0$. Thus, the dimension of $H^0(X^F, \mathcal{C})$ can be computed by Riemann-Roch. Finally, the

description of $H^1(X^F, \mathcal{K})^\vee$ follows immediately from applying duality to the computation in (i). \circ

Let us assume from now on that $X^{\log} \rightarrow S^{\log}$ is an r -pointed stable curve of genus g ; (P, ∇_P) is crys-stable; and \mathcal{W} is mild. Thus, we have $\mathrm{Ad}^q(P) \subseteq \mathrm{Ad}(P)$, and the de Rham cohomology $H_{\mathrm{DR}}^i(X, \mathcal{W})$ is defined by considering the complex $\mathrm{Ad}(P) \rightarrow \mathrm{Ad}^q(P) \otimes \omega_{X^{\log}/S^{\log}}$.

Next, we would like to consider the image \mathcal{I} of $H_{\mathrm{DR}}^1(X, \mathrm{Ad}(P))$ in $H_{\mathrm{DR}}^1(X, \mathcal{W})$. Let $\mathcal{H} \subseteq \mathrm{Ad}(P)$ be the kernel of the surjection $\mathrm{Ad}(P) \rightarrow \mathcal{W}$ induced by \mathcal{P}^\vee . Then \mathcal{H} inherits a connection $\nabla_{\mathcal{H}}$ from $\mathrm{Ad}(P)$, and, moreover, fits into an exact sequence

$$0 \rightarrow \mathcal{W}^\vee \rightarrow \mathcal{H} \rightarrow \mathcal{O}_X \rightarrow 0$$

which is stabilized by $\nabla_{\mathcal{H}}$. At any rate, the cokernel of $H_{\mathrm{DR}}^1(X, \mathrm{Ad}(P)) \rightarrow H_{\mathrm{DR}}^1(X, \mathcal{W})$ injects into $H_{\mathrm{DR}}^2(X^{\log}, \mathcal{H})$, i.e., the second hypercohomology group of the complex $\mathcal{H} \rightarrow (\mathcal{H} \cap \mathrm{Ad}^q(P)) \otimes \omega_{X^{\log}/S^{\log}}$ defined by $\nabla_{\mathcal{H}}$. On the other hand, $H_{\mathrm{DR}}^2(X^{\log}, \mathcal{H})$ is easily seen to be dual to $H^0(X, \mathcal{O}_X(-M_X^{\mathrm{non}}))^F$ (since the p -curvature of \mathcal{H}^\vee is generically nonzero). Moreover, one sees easily (indeed, this is another special case of the general principle “the derivative of the p -curvature (or, in this case, the p -curvature in the case of an affine structure group) is minus the Cartier operator” – cf., Lemma 2.7 above; the proof of Theorem 2.13 of [Mzk1], Chapter II) that the morphism

$$H^1(X^F, \mathcal{K}) \hookrightarrow H_{\mathrm{DR}}^1(X, \mathcal{W}) \rightarrow H_{\mathrm{DR}}^2(X^{\log}, \mathcal{H}) =$$

$$(H^0(X, \mathcal{O}_X(-M_X^{\mathrm{non}}))^F)^\vee$$

is simply (minus) the surjection

$$\zeta : H^1(X, \mathcal{K}) = H^0(Z^F, \mathcal{O}_{Z^F}(D_Z - M_Z^F))^\vee \rightarrow$$

$$(H^0(X, \mathcal{O}_X(-M_X^{\mathrm{non}}))^F)^\vee$$

(where the second morphism is the dual to the natural inclusion). Thus, we obtain the following:

Proposition 3.5. *We have an exact sequence*

$$0 \rightarrow \mathrm{Ker}(\zeta) \rightarrow \mathcal{I} \rightarrow$$

$$H^0(X^F, \mathcal{C}) \rightarrow 0$$

where \mathcal{I} is the image of $H_{\text{DR}}^1(X, \text{Ad}(P))$ in $H_{\text{DR}}^1(X, \mathcal{W})$.

Next, we would like to consider the first hypercohomology module $H_{\text{DR}}^1(X^{\log}, \mathcal{O}_X)$ of the de Rham complex

$$\mathcal{O}_X \xrightarrow{d} \omega_{X^{\log}/S^{\log}}(-M_X^{\text{ac}})$$

on X . By a similar spectral sequence to the one just considered above, we have an exact sequence

$$0 \rightarrow H^1(X, \mathcal{O}_X)^F \rightarrow H_{\text{DR}}^1(X^{\log}, \mathcal{O}_X) \rightarrow$$

$$H^0(X, \omega_{X^{\log}/S^{\log}}(-M_X^{\text{ac}}))^F \rightarrow 0$$

Now let us recall that we have a horizontal exact sequence $0 \rightarrow \mathcal{W}^\vee \rightarrow \mathcal{H} \rightarrow \mathcal{O}_X \rightarrow 0$. The cokernel of the natural morphism $H_{\text{DR}}^1(X^{\log}, \mathcal{H}) \rightarrow H_{\text{DR}}^1(X^{\log}, \mathcal{O}_X)$ thus injects into

$$\begin{aligned} H_{\text{DR}}^2(X^{\log}, \mathcal{W}^\vee) &\cong H^0(X^F, \mathcal{K}(-M_X^F|_{Z^F}))^\vee \\ &\cong H^0(Z^F, \omega_{Z^F/k}(-D_Z + M_Z^F - M_X^F|_{Z^F}))^\vee \\ &\cong H^1(Z^F, \mathcal{O}_{Z^F}(D_Z + M_X^F|_{Z^F} - M_Z^F)) \end{aligned}$$

(Here we use duality and Proposition 3.4.) On the other hand, it is easy to see (yet another special case of the general principle “the p -curvature in the case of an affine structure group is minus the Cartier operator” – cf., Lemma 2.7 above; the proof of Theorem 2.13 of [Mzk1], Chapter II) that the composite

$$\begin{aligned} H^1(X, \mathcal{O}_X)^F &\subseteq \\ H_{\text{DR}}^1(X^{\log}, \mathcal{O}_X) &\rightarrow H_{\text{DR}}^2(X^{\log}, \mathcal{W}^\vee) \\ &= H^1(Z^F, \mathcal{O}_{Z^F}(D_Z + (M_X|_Z - M_Z)^F)) \end{aligned}$$

is simply (minus) the natural morphism induced on H^1 's by the inclusion of sheaves $\mathcal{O}_{X^F} \hookrightarrow (\mathcal{O}_Z(D_Z + M_X|_Z - M_Z))^F$, hence is *surjective*. Thus, we obtain the following:

Proposition 3.6. *Let \mathcal{J} be the image of $H_{\text{DR}}^1(X^{\log}, \mathcal{H})$ in $H_{\text{DR}}^1(X^{\log}, \mathcal{O}_X)$. Then the restriction of the natural surjection*

$$H_{\text{DR}}^1(X^{\log}, \mathcal{O}_X) \rightarrow H^0(X, \omega_{X^{\log}/S^{\log}}(-M_X^{\text{ac}}))^F$$

considered above to \mathcal{J} remains surjective.

Now let us note that we have a natural inclusion

$$(\omega_{Z^{\log}/S^{\log}})^F(-D_Z) \subseteq \omega_{X^{\log}/S^{\log}}^F(-M_X^{\text{ac}})^F$$

of sheaves on X . Let us denote the cokernel of this natural inclusion by \mathcal{G} . Let \mathcal{F} be the quotient of $H^0(X, \omega_{X^{\log}/S^{\log}}(-M_X^{\text{ac}}))^F$ by

$$H^0(Z^F, (\omega_{Z^{\log}/S^{\log}})^F(-D_Z)) = H^0(X^F, \mathcal{K})$$

Thus, we have a surjection

$$H_{\text{DR}}^1(X^{\log}, \mathcal{O}_X) \rightarrow \mathcal{F}$$

Proposition 3.7. *We have: $\dim_k(\text{Ker}(\zeta)) + \dim_k(\mathcal{F}) = \deg(D_Z) - \deg(M_Z^{\text{ac}}) + N_X^{\text{ac}}$.*

Proof. Consider the long exact cohomology sequence associated to

$$0 \rightarrow (\omega_{Z^{\log}/S^{\log}})^F(-D_Z) \rightarrow (\omega_{X^{\log}/S^{\log}}(-M_X^{\text{ac}}))^F \rightarrow \mathcal{G} \rightarrow 0$$

and observe that $\deg(\mathcal{G}) = \deg(D_Z) - \deg(M_Z^{\text{ac}}) + N_X^{\text{ac}}$. \circ

Finally, before continuing, we would like to give one more technical lemma which will be of use to us in the following §. We continue to let k be an algebraically closed field of characteristic p . Let A be a local artinian k -algebra with residue field k . Let $S \stackrel{\text{def}}{=} \text{Spec}(A)$. Let $I \subseteq A$ be an ideal of length (as an A -module) 1. Let us assume that S is equipped with some fine log structure; denote the resulting log scheme by S^{\log} . Let $X^{\log} \rightarrow S^{\log}$ be an r -pointed stable log-curve of genus g , where $2g - 2 + r \geq 1$. Let us employ a subscript “ I ” (respectively, “ 0 ”) to denote the reductions of objects over A modulo I (respectively, the maximal ideal of A). Thus, $S_I = \text{Spec}(A/I)$.

Lemma 3.8. *Suppose that we are given a line bundle with connection $(\mathcal{L}_I, \nabla_{\mathcal{L}_I})$ on X_I^{\log} such that the monodromy operators defined by $\nabla_{\mathcal{L}_I}$ at the marked points of X_I^{\log} are given by multiplication by elements of \mathbf{F}_p . Then this line bundle with connection lifts to a line bundle with connection $(\mathcal{L}, \nabla_{\mathcal{L}})$ on X^{\log} whose monodromy operators at the marked points are given by multiplication by elements of \mathbf{F}_p .*

Proof. First, observe that there exists a power series ring B over k (i.e., a ring isomorphic to $k[[x_1, \dots, x_N]]$, where x_1, \dots, x_N are indeterminates) that surjects onto A . Thus, we shall think of A as a quotient of B . Let $T \stackrel{\text{def}}{=} \text{Spec}(B)$. Without loss of generality (that is to say, up to possibly replacing the log structure on S by the log structure obtained by pulling back the log structure on $\overline{\mathcal{M}}_{g,r}^{\log}$ by the classifying morphism),

we may assume that the log-curve $X^{\log} \rightarrow S^{\log}$ arises as the restriction to S of a log-curve $X_T^{\log} \rightarrow T^{\log}$ such that $X_T \rightarrow T$ is *smooth* over the generic point of T . Also, since the obstruction to lifting a line bundle (without a connection) lies in H^2 (in the Zariski topology) of the structure sheaf, and such an H^2 necessarily vanishes in the case of a curve, it follows that we may assume that \mathcal{L}_I lifts to a line bundle \mathcal{L}_T on X_T . Let us write ω_{X_T} for the relative dualizing sheaf of the morphism $X_T \rightarrow T$. Now the obstruction to putting a connection on \mathcal{L}_T (whose monodromy operators are as stipulated in the lemma) defines a class $c(\mathcal{L}_T) \in H^1(X_T, \omega_{X_T})$. Next, recall that the “trace map” gives an isomorphism $H^1(X_T, \omega_{X_T}) \cong B$. Moreover, it follows from the theory of the trace map that

The image of $c(\mathcal{L}_T)$ relative to this isomorphism is the degree $\deg(\mathcal{L}_T) \in B$ of the line bundle \mathcal{L}_T , regarded as an element of $\mathbf{F}_p \subseteq B$.

Indeed, since B (unlike A) is a domain, it suffices to verify the italicized assertion over the quotient field of B . But, in the case of a smooth curve over a field, the italicized assertion follows from the discussion at the end of [Harts2], Chapter III, §7 (cf. also Exercises 7.2, 7.4 of that §). This completes the verification of the italicized assertion above.

In particular, it follows that *whether or not $c(\mathcal{L}_T)$ is zero may be checked after restriction to S_I* (where it is already known to be zero). Thus, we conclude that $c(\mathcal{L}_T) = 0$, i.e., that \mathcal{L}_T admits a connection with the desired properties. Moreover, since the connections on \mathcal{L}_T (with prescribed monodromy) form a torsor over $H^0(X_T, \omega_{X_T})$, and the natural restriction morphism

$$H^0(X_T, \omega_{X_T}) \rightarrow H^0(X, \omega_X)$$

is clearly surjective, it thus follows that \mathcal{L}_T admits a connection (with the desired monodromy) that lifts $\nabla_{\mathcal{L}_I}$. This completes the proof of the lemma. \circ

§3.3. Deformation Theory

Let us maintain the notations introduced at the end of the preceding §: Thus, $f^{\log} : X^{\log} \rightarrow S^{\log}$ is an r -pointed stable log-curve of genus g , where $2g - 2 + r \geq 1$; S^{\log} is a fine noetherian log scheme of characteristic $p > 2$; and $S = \operatorname{Spec}(A)$, where A is an artinian local k -algebra with residue field k , and k is algebraically closed. Write $S_0 \subseteq S$ for $\operatorname{Spec}(k) \subseteq \operatorname{Spec}(A)$. Let $I \subseteq A$ be an ideal of length 1. Thus, as an A -module, $I \cong k$. Let us employ a subscript “ I ” (respectively, “ 0 ”) to denote the reductions of objects over A modulo I (respectively, the maximal ideal of A). Thus, $S_I = \operatorname{Spec}(A/I)$. Finally, we suppose that we are given a \mathbf{P}^1 -bundle with connection (P, ∇_P) on X^{\log} .

Now fix a collection ρ of r radii in \mathbf{F}_p . Let us assume that $(P, \nabla_P)_I$ is a mildly spiked torally indigenous bundle of radii ρ on X_I^{\log} , and let us consider the obstruction to lifting it to a mildly spiked nilpotent torally crystal stable bundle of radii ρ on X^{\log} . Let $V_I \subseteq X_I$ be the S_I -flat divisor which is the spiked locus of $(P, \nabla_P)_I$. Thus, we have a horizontal (square) nilpotent surjection

$$\kappa_I : \mathrm{Ad}(P_I) \rightarrow \mathcal{W}_I \stackrel{\mathrm{def}}{=} (\Phi_{X^{\log}/S^{\log}}^* \omega_{X^{\log}/S^{\log}})_I(-V_I)$$

Write \mathcal{W}_0 for $(\mathcal{W}_I)_0$. Then, by the general nonsense of deformation theory, the obstruction to lifting this surjection to a horizontal (square) nilpotent surjection of $\mathrm{Ad}(P)$ (onto some line bundle) forms an element $\eta \in I \otimes_k H_{\mathrm{DR}}^1(X_0, \mathcal{W}_0)$. I claim that this element η vanishes under the connecting homomorphism

$$I \otimes_k H_{\mathrm{DR}}^1(X_0, \mathcal{W}_0) \rightarrow I \otimes_k H_{\mathrm{DR}}^2(X_0^{\log}, \mathcal{O}_{X_0})$$

studied in the preceding subsection (cf. the discussion preceding Proposition 3.5). (Here, the “ H_{DR}^2 ” is the second hypercohomology module of the complex appearing in the discussion following Proposition 3.5.) Indeed, to see this, it suffices to observe the following: First of all, by Lemma 3.8, the line bundle \mathcal{W}_I equipped with its natural connection lifts to a line bundle with connection $(\mathcal{L}, \nabla_{\mathcal{L}})$ on X^{\log} whose monodromy at each marked point is the same as that of \mathcal{W}_I . Let $\mathcal{B}_I = \mathrm{Ad}(P_I)/\mathrm{Im}((\kappa_I)^\vee)$. Then it follows from the general nonsense of deformation theory that the obstruction to lifting κ_I to a horizontal (square) nilpotent surjection of $\mathrm{Ad}(P)$ onto the particular line bundle with connection $(\mathcal{L}, \nabla_{\mathcal{L}})$ forms an element $\eta_{\mathcal{B}} \in I \otimes_k H_{\mathrm{DR}}^1(X_0, \mathcal{B}_0)$ which maps to η via the projection $\mathcal{B}_0 \rightarrow \mathcal{W}_0$. (Here, the de Rham cohomology module $H_{\mathrm{DR}}^1(X_0, \mathcal{B}_0)$ is defined by means of the quotient (induced by the surjection $\mathrm{Ad}(P_0) \rightarrow \mathcal{B}_0$) of the complex defining the de Rham cohomology of $\mathrm{Ad}(P_0)$.) The existence of $\eta_{\mathcal{B}}$ thus completes the proof of the claim that η maps to zero under the above connecting homomorphism.

Thus, it follows (from the discussion preceding Proposition 3.5) that η lies in $I \otimes_k \mathcal{I}_0$, i.e., the image of $I \otimes_k H_{\mathrm{DR}}^1(X_0, \mathrm{Ad}(P_0))$ in $I \otimes_k H_{\mathrm{DR}}^1(X_0, \mathcal{W}_0)$. But this means that by suitably modifying the deformation (P, ∇_P) of $(P, \nabla_P)_I$, we may always assume that there exists a horizontal (square) nilpotent surjection

$$\kappa' : \mathrm{Ad}(P) \rightarrow \mathcal{W}'$$

whose reduction modulo I is equal to κ_I . Note that the choice of such a κ' is not necessarily unique – indeed, such κ' form a torsor over $H_{\mathrm{DR}}^0(X_0, \mathcal{W}_0) = H^0(X_0^F, \mathcal{K}_0)$, where $\mathcal{K}_0 \subseteq \mathcal{W}_0$ is the subsheaf of horizontal sections.

Let $\mathcal{B}' = \mathrm{Ad}(P)/\mathrm{Im}((\kappa')^\vee)$. Let $\psi' : \mathrm{Ad}(P) \rightarrow \mathcal{B}'$ be the natural surjection. Note that we have a horizontal exact sequence $0 \rightarrow \mathcal{O}_X \rightarrow \mathcal{B}' \rightarrow \mathcal{W}' \rightarrow 0$. The existence of such a horizontal exact sequence

implies that the image of the p -curvature of \mathcal{B}' , hence of (P, ∇_P) , lies in $\text{Ker}(\kappa') \subseteq \text{Ad}(P)$. Thus, the composite of the p -curvature $\mathcal{P} : \Phi_X^* \tau_{X^{\log}/S^{\log}} \rightarrow \text{Ad}(P)$ with κ' must vanish. Unfortunately, however, the composite of the p -curvature $\mathcal{P} : \Phi_X^* \tau_{X^{\log}/S^{\log}} \rightarrow \text{Ad}(P)$ with ψ' might not vanish. This composite will, however, vanish modulo I . Thus, it defines a horizontal morphism $\Phi_X^*(\tau_{X^{\log}/S^{\log}})_0 \rightarrow \mathcal{O}_{X_0}$, i.e., an element ξ of $H^0(X_0, (\omega_{X^{\log}/S^{\log}})_0(-M_{X_0}^{\text{ac}}))^F$. In fact, there is an ambiguity in the choice of κ' , so, taking this ambiguity into account, we obtain an element $\delta \in \mathcal{F}_0$ (which is the image of ξ under the surjection $H^0(X_0, (\omega_{X^{\log}/S^{\log}})_0(M_{X_0}^{\text{ac}}))^F \rightarrow \mathcal{F}_0$). On the other hand, since $H_{\text{DR}}^1(X_0^{\log}, \text{Ker}(\kappa'_0)) \rightarrow H^0(X_0, (\omega_{X^{\log}/S^{\log}})_0(-M_{X_0}^{\text{ac}}))^F \rightarrow \mathcal{F}_0$ is surjective (Proposition 3.6), it thus follows that we may choose a deformation (P, ∇_P) of $(P, \nabla_P)_I$ such that $\delta = 0$.

Now let us interpret what this means. We have just seen that we may choose a deformation (P, ∇_P) of $(P, \nabla_P)_I$ such that the p -curvature \mathcal{P} of (P, ∇_P) factors through a locally split horizontal injection

$$(\mathcal{W})^\vee \hookrightarrow \text{Ad}(P)$$

whose image is square nilpotent. But this means that $V_{\mathcal{P}}$ is a *divisor*, which is necessarily S -flat. Thus, (after adjusting the deformation X^{\log} of X_I^{\log} so that (P, ∇_P) becomes torally indigenous) we see that (P, ∇_P) is a *mildly spiked torally indigenous bundle*. In other words, we have proven the smoothness part of the following:

Theorem 3.9. *The stack $\overline{\mathcal{N}}_{g,r}^\rho[d]$ of mildly spiked torally indigenous bundles of strength d and radii ρ is smooth over \mathbf{F}_p of dimension $3g - 3 + r$ (if it is nonempty). In particular, the natural morphism $\overline{\mathcal{N}}_{g,r}^\rho[d] \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ is flat and quasi-finite.*

Proof. It remains only to compute the dimension. But, from the above discussion, one sees that this dimension is equal to

$$2(3g - 3 + r) - \dim_k(\mathcal{I}_0) - \dim_k(\mathcal{F}_0) = 3g - 3 + r$$

(by Propositions 3.4, 3.5, 3.6, and 3.7). \circ

It remains to consider what happens when X_0^{\log} is *totally degenerate*. Thus, assume that X_0^{\log} is totally degenerate. Let us denote by $\mathcal{M} \subseteq H_{\text{DR}}^1(X_0, \text{Ad}(P_0))$ the infinitesimal deformations that are mildly spiked. By Theorem 3.9, $\dim_k(\mathcal{M}) = 3g - 3 + r$. Let $Q = H^0(X, (\omega_{X^{\log}/S^{\log}})^{\otimes 2}(-M_X))_0 \subseteq H_{\text{DR}}^1(X_0, \text{Ad}(P_0))$. I claim that $Q \cap \mathcal{M} = 0$. Indeed, to see this, it suffices to prove that mildly spiked torally indigenous bundles on totally degenerate curves admit no nontrivial deformations. But observe that the radii at the nodes of such a bundle can be determined from the eigenvalues of the monodromy of the line bundle $(\Phi_X^* \omega_{X^{\log}/S^{\log}})(-\mathcal{V}_{\mathcal{P}})$. Moreover, these eigenvalues are clearly $\in \mathbf{F}_p$. Thus, the radii at the

nodes also $\in \mathbf{F}_p$, so by Chapter I, Theorem 4.4, we conclude that such bundles have no nontrivial deformations. This completes the proof of the claim. Thus, we have proven the following result:

Proposition 3.10. *The natural morphism $\overline{\mathcal{N}}_{g,r}^\rho[d] \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ is étale over points corresponding to totally degenerate curves.*

Chapter III: VF-Patterns

§0. Introduction

In this Chapter, we set up the *uniformization theory* (to be discussed in the rest of the book) *modulo* p . Ultimately, the uniformization theory will be based on considering moduli of indigenous bundles with various kinds of Frobenius actions. The problem is that in this book, we wish to allow for the possibility of many different types of Frobenius actions (on an indigenous bundle), not just the simplest type, which is studied in [Mzk1]. The *combinatorial data* which describes what kind of Frobenius action one is dealing with is called a *VF-pattern*. The goal of this Chapter, then, is to construct moduli stacks – called *VF-stacks* – of “data modulo p ” for each VF-pattern. We then show that these moduli stacks have various *smoothness* and *affineness* properties (Proposition 1.7 and Theorem 2.10). Finally, we note that the various properties of these moduli stacks allow one to give a new proof of the *connectedness* of $\overline{\mathcal{M}}_{g,r}$. As usual, we assume throughout that p is an *odd* prime.

§1. The Moduli Stack Associated to a VF-Pattern

§1.1. Definition of a VF-pattern

Fix nonnegative integers g and r such that $2g - 2 + r \geq 1$. Let $\chi = \frac{1}{2}(2g - 2 + r)$. Let ϖ be a positive integer. Let $\mathfrak{L}\mathfrak{e}\mathfrak{v}$ be the finite set consisting of the half-integers between (and including) 0 and χ . We shall call $\mathfrak{L}\mathfrak{e}\mathfrak{v}$ the *set of levels*, and we shall refer to an element of $\mathfrak{L}\mathfrak{e}\mathfrak{v}$ as a *level*. Let $\Pi : \mathbf{Z} \rightarrow \mathfrak{L}\mathfrak{e}\mathfrak{v}$ be a map of sets.

Definition 1.1. We shall call a pair (Π, ϖ) a *VF-pattern of period* ϖ if the following conditions are satisfied: $\Pi(n + \varpi) = \Pi(n)$ for all $n \in \mathbf{Z}$; $\Pi(0) = \chi$; for any two

$i, j \in \mathbf{Z}$, $\Pi(i) - \Pi(j) \in \mathbf{Z}$. If Π maps every integer to χ (and ϖ is arbitrary), we shall call Π the *pre-home VF-pattern of period ϖ* . If $\varpi = 1$, we shall call Π the *home VF-pattern*. Finally, if $\Pi(\mathbf{Z}) \subseteq \{0, \chi\}$, then we shall call Π *binary*.

Note in particular that a VF-pattern of period ϖ automatically also defines a VF-pattern of period $n \cdot \varpi$ for any positive integer n . Often it is useful to think of a VF-pattern as a sequence of levels indexed by the integers (not shown):

$$\cdots \quad \frac{5}{2} \quad \frac{1}{2} \quad \chi \quad \frac{7}{2} \quad \frac{3}{2} \quad \cdots$$

We shall refer to a term in this sequence at a *link of the VF-pattern Π* . If $a \in \mathbf{Z}$ is such that $\Pi(a) = \chi$, we shall call a *indigenous for Π* . We shall refer to $a \in \mathbf{Z}$ such that $\Pi(a) \neq 0$ (respectively, $\Pi(a) = 0$) as *active (respectively, dormant) for Π* . If $a, b \in \mathbf{Z}$ are such that $a < b$; $\Pi(a)$ and $\Pi(b)$ are active (respectively, indigenous); and $\Pi(c)$ is dormant (respectively, not indigenous) for all c such that $a < c < b$, then we shall call such (a, b) Π -*adjacent* (respectively, Π -*ind-adjacent*).

§1.2. Construction of Link Stacks

Let i and j be levels (i.e., elements of $\mathfrak{L}\mathfrak{ev}$). Let $(\overline{\mathcal{Y}}_{g,r}^i)_{\mathbf{F}_p}$ be the moduli stack over \mathbf{F}_p of r -pointed stable curves of genus g equipped with a crys-stable bundle of level i . In this subsection, we construct a certain substack

$$\mathcal{R}_{g,r}^{i,j} \subseteq (\overline{\mathcal{Y}}_{g,r}^i)_{\mathbf{F}_p}$$

which we shall call the *link stack between levels i and j* . We define $\mathcal{R}_{g,r}^{i,j}$ as follows:

Definition 1.2. Let S be a \mathbf{F}_p -scheme. Then we take $\mathcal{R}_{g,r}^{i,j}(S)$ to be the full subcategory of $(\overline{\mathcal{Y}}_{g,r}^i)_{\mathbf{F}_p}(S)$ whose objects are the pairs $(f^{\log} : X^{\log} \rightarrow S^{\log}; (\pi : P \rightarrow X, \nabla_P))$ (where $(P \rightarrow X, \nabla_P)$ is a crys-stable bundle of level i on the log-curve $f^{\log} : X^{\log} \rightarrow S^{\log}$) such that the *p -curvature*

$$\mathcal{P} : \Phi_X^*(\tau_{X^{\log}/S^{\log}}) \rightarrow \mathrm{Ad}(P)$$

(where $\Phi_X : X \rightarrow X$ is the absolute Frobenius morphism of X) of $(P \rightarrow X, \nabla_P)$ satisfies the following:

- (1) $\kappa(\mathcal{P}, \mathcal{P}) : \Phi_X^*(\tau_{X^{\log}/S^{\log}})^{\otimes 2} \rightarrow \mathcal{O}_X$ is zero (i.e., the image of \mathcal{P} is square nilpotent);

- (2) if $j = 0$, then $f : X \rightarrow S$ is *smooth*, and \mathcal{P} is identically zero;
- (3) if $j > 0$, then the zero locus (as a closed subscheme of X) of \mathcal{P}^\vee (the dual of \mathcal{P}) is the schematic inverse image via the relative Frobenius morphism $\Phi_{X/S} : X \rightarrow X^F \stackrel{\text{def}}{=} X \times_{S, \Phi_S} S$ of a “ j -balanced” divisor $D_{\mathcal{P}} \subseteq X^F$ (cf. Definition 3.7 of Chapter I for a definition of this notion). We shall refer to $D_{\mathcal{P}} \subseteq X^F$ as the *p -curvature locus* of (P, ∇_P) .

We shall refer to such a $(P \rightarrow X, \nabla_P)$ as an (i, j) -link bundle.

Proposition 1.3. *The morphism $\mathcal{R}_{g,r}^{i,j} \hookrightarrow (\overline{\mathcal{Y}}_{g,r}^i)_{\mathbb{F}_p}$ is representable by a subalgebraic stack of $(\overline{\mathcal{Y}}_{g,r}^i)_{\mathbb{F}_p}$.*

Proof. Condition (1) of Definition 1.2 clearly defines a closed subalgebraic stack of $(\overline{\mathcal{Y}}_{g,r}^i)_{\mathbb{F}_p}$. Similarly, condition (2) of Definition 1.2 defines a locally closed sub-algebraic stack of $(\overline{\mathcal{Y}}_{g,r}^i)_{\mathbb{F}_p}$. Thus, it suffices to consider condition (3) of Definition 1.2 in the case $j > 0$. Since the stack of stable curves equipped with a j -balanced divisor (i.e., the stack $\overline{\mathcal{D}}_{g,r}^j$ of Definition 3.7 of Chapter I) is algebraic, the representability of $\mathcal{R}_{g,r}^{i,j}$ follows from the fact that given two S -flat divisors $D_1, D_2 \subseteq X$, there exists a natural closed subscheme of S over which D_1 and D_2 coincide. \square

Now let $(P \rightarrow X; \nabla_P)$ be an (i, j) -link bundle on $f^{\log} : X^{\log} \rightarrow S^{\log}$. Let $\Phi_{X/S} : X \rightarrow X^F$ be the relative Frobenius morphism over S . Let us examine the case $j = 0$. In this case, the p -curvature of the vector bundle $\text{Ad}(P)$ is identically zero. If X^{\log} had a marked point, the monodromy at that marked point would be nonzero and square-nilpotent; thus, the p -curvature at that marked point would be nonzero (cf. the discussion preceding Proposition 1.3 of Chapter II). It thus follows that X^{\log} cannot have any marked points. Moreover, by assumption, $f : X \rightarrow S$ is *smooth*. It thus follows from the basic theory of p -curvature that there exists a \mathbf{P}^1 -bundle $Q \rightarrow X^F$ over X^F such that we have an isomorphism $\Phi_{X/S}^* Q \cong P$ with respect to which the sections of $P \rightarrow X$ that arise from sections of $Q \rightarrow X^F$ are horizontal for ∇_P . Moreover, it follows immediately from the fact that $(P \rightarrow X; \nabla_P)$ is crys-stable that for *any* connection ∇_Q on $Q \rightarrow X^F$, the pair $(Q \rightarrow X^F; \nabla_Q)$ forms a crys-stable bundle of level 0 on X^F . We summarize this as follows:

Proposition 1.4. *If there exists an $(i, 0)$ -link bundle on X^{\log} , then X^{\log} has no marked points, and any $(i, 0)$ -link bundle can be obtained (Zariski-locally on S) as the pull-back via $\Phi_{X/S}$ of a crys-stable bundle of level 0 on X^F .*

Next, let us examine the case $j > 0$. Let $D_{\mathcal{P}} \subseteq X^F$ be the j -balanced divisor which is the p -curvature locus of (P, ∇_P) . It thus follows that \mathcal{P} extends to a *locally split* morphism

$$\mathcal{P} : \Phi_{X/S}^* \{(\tau_{X^{\log}/S^{\log}})^F(D_{\mathcal{P}})\} \rightarrow \mathrm{Ad}(P)$$

Let $\mathcal{K} \subseteq \mathrm{Ad}(P)$ be the kernel of \mathcal{P}^\vee . The logarithmic connection on $\mathrm{Ad}(P)$ induces a natural logarithmic connection $\nabla_{\mathcal{K}}$ on \mathcal{K} such that $(P \rightarrow X; \nabla_P)$ is naturally isomorphic to $\mathbf{P}(\mathcal{K}, \nabla_{\mathcal{K}})$. Moreover, the image of \mathcal{P} lies in \mathcal{K} , and the quotient of \mathcal{K} by the image of \mathcal{P} is naturally isomorphic to \mathcal{O}_X . In other words, we have an exact sequence of vector bundles with logarithmic connections on X^{\log}

$$0 \rightarrow \Phi_{X/S}^* \{(\tau_{X^{\log}/S^{\log}})^F(D_{\mathcal{P}})\} \rightarrow \mathcal{K} \rightarrow \mathcal{O}_X \rightarrow 0$$

Let $\mathcal{T} = \Phi_{X/S}^* \{(\tau_{X^{\log}/S^{\log}})^F(D_{\mathcal{P}})\}$; we equip \mathcal{T} with the connection $\nabla_{\mathcal{T}}$ for which sections of $\Phi_{X/S}^{-1} \{(\tau_{X^{\log}/S^{\log}})^F(D_{\mathcal{P}})\}$ are horizontal. Then the extension class corresponding to the above exact sequence defines a section s_P of

$$\mathbf{R}^1(f_{\mathrm{DR}})_* \mathcal{T}$$

i.e., the first de Rham cohomology module of $(\mathcal{T}, \nabla_{\mathcal{T}})$ with respect to $f^{\log} : X^{\log} \rightarrow S^{\log}$ over S . That is to say, this cohomology module is the first hypercohomology module of the complex

$$\mathcal{T} \xrightarrow{\nabla_{\mathcal{T}}} \mathcal{T} \otimes_{\mathcal{O}_X} \omega_{X/S}^{\log}$$

of sheaves on X . Computing this cohomology module by means of the so-called “conjugate spectral sequence” (cf., e.g., the discussion of §2.2 of Chapter II, especially Corollary 2.6; §8.25 of [BO]) gives an exact sequence

$$0 \rightarrow \Theta_{D_{\mathcal{P}}} \rightarrow \mathbf{R}^1(f_{\mathrm{DR}})_* \mathcal{T} \rightarrow f_*(\mathcal{O}_{X^F}(D_{\mathcal{P}})) \rightarrow 0$$

where we denote by $\Theta_{D_{\mathcal{P}}}$ the vector bundle $\mathbf{R}^1 f_*^F \{(\tau_{X^{\log}/S^{\log}})^F(D_{\mathcal{P}})\}$. (Here $f^F : X^F \rightarrow S$ is the result of base-changing f by $\Phi_S : S \rightarrow S$, the absolute Frobenius morphism on S .) Let us denote by

$$\mathbf{R}^1(f_{\mathrm{DR}})'_* \mathcal{T} \subseteq \mathbf{R}^1(f_{\mathrm{DR}})_* \mathcal{T}$$

the subsheaf obtained by pulling back the preceding exact sequence by the natural inclusion

$$\mathcal{O}_S \hookrightarrow f_* \mathcal{O}_{X^F} \hookrightarrow f_*(\mathcal{O}_{X^F}(D_{\mathcal{P}}))$$

Thus, we have an exact sequence

$$0 \rightarrow \Theta_{D_{\mathcal{P}}} \rightarrow \mathbf{R}^1(f_{\mathrm{DR}})'_* \mathcal{T} \rightarrow \mathcal{O}_S \rightarrow 0$$

Let

$$\mathcal{A} \rightarrow S$$

denote the (geometric) $\Theta_{D_{\mathcal{P}}}$ -torsor of splittings of the above exact sequence. Now just as in [Mzk1], Chapter II, §1, one sees immediately that $s_{\mathcal{P}}$ lies in the primed de Rham cohomology group, and that the image of $s_{\mathcal{P}}$ in \mathcal{O}_S is *invertible*. Thus, $s_{\mathcal{P}}$ defines a section $a_{\mathcal{P}}$ of $\mathcal{A} \rightarrow S$ over S .

Conversely, given any section $a : S \rightarrow \mathcal{A}$ of $\mathcal{A} \rightarrow S$, there exists a section s of $\mathbf{R}^1(f_{\mathrm{DR}})'_* \mathcal{T}$ which gives rise to a . Thus, s defines an extension of vector bundles with logarithmic connections on X^{\log}

$$0 \rightarrow \Phi_{X/S}^* \{(\tau_{X^{\log}/S^{\log}})^F(D_{\mathcal{P}})\} \rightarrow \mathcal{F}_a \rightarrow \mathcal{O}_X \rightarrow 0$$

Let us denote by $(P_a \rightarrow X, \nabla_{P_a})$ the \mathbf{P}^1 -bundle with connection defined by $(\mathcal{F}_a, \nabla_{\mathcal{F}_a})$. We claim that $(P_a \rightarrow X, \nabla_{P_a})$ is an (i, j) -link bundle (with p -curvature locus equal to $D_{\mathcal{P}}$). Indeed, first one observes – by the principle that the p -curvature for an “affine structure group” (such as \mathcal{T}) is just the Cartier operator, i.e., the projection $\mathbf{R}^1(f_{\mathrm{DR}})'_* \mathcal{T} \rightarrow f_*(\mathcal{O}_{X^F}(D_{\mathcal{P}}))$ of the exact sequence considered above (cf., e.g., [Mzk1], Chapter II: Proposition 1.4, the proof of Theorem 2.13) – that the p -curvature of $(P_a \rightarrow X, \nabla_{P_a})$ defines a locally split morphism

$$\mathcal{P}_a : \Phi_{X/S}^* \{(\tau_{X^{\log}/S^{\log}})^F(D_{\mathcal{P}})\} \rightarrow \mathrm{Ad}(P_a)$$

whose image is square-nilpotent. Now one can read off the fact that $(P_a \rightarrow X, \nabla_{P_a})$ is crys-stable by just looking at \mathcal{P}_a :

- (1) condition (1) of Definition 1.2 of Chapter I follows from the construction of \mathcal{F}_a as an extension of line bundles with connection whose p -curvature is zero;
- (2) conditions (2) and (4) of Definition 1.2 of Chapter I follow from the fact that $D_{\mathcal{P}}$ is j -balanced, hence avoids nodes and marked points;
- (3) condition (3) of Definition 1.2 of Chapter I follows from the fact that any offending $\alpha : \mathcal{L} \hookrightarrow \mathrm{Ad}(P_a)|_Y$

would have to be generically distinct from the image of \mathcal{P}_a but contained in $\mathcal{K}_a \stackrel{\text{def}}{=} \text{Ker}(\mathcal{P}_a^\vee) \subseteq \text{Ad}(\mathcal{P}_a)$ (since $\text{Ad}(\mathcal{P}_a)/\text{Im}(\mathcal{P}_a)$ has nonzero p -curvature), but this would imply that $(\mathcal{K}_a, \nabla_{\mathcal{K}_a}) \cong (\mathcal{F}_a, \nabla_{\mathcal{F}_a})$ has p -curvature zero on Y , which is absurd.

We summarize this as follows:

Proposition 1.5. *Let j be a nonzero level. Let $D \subseteq X^F$ be a j -balanced divisor. Let $\mathcal{T} \stackrel{\text{def}}{=}} \Phi_{X/S}^* \{(\tau_{X^{\log}/S^{\log}})^F(D)\}$; we equip \mathcal{T} with its natural connection. Let*

$$\mathcal{A} \rightarrow S$$

denote the (geometric) Θ_D -torsor of splittings of the exact sequence

$$0 \rightarrow \Theta_D \rightarrow \mathbf{R}^1(f_{\text{DR}})_* \mathcal{T} \rightarrow \mathcal{O}_S \rightarrow 0$$

where $\Theta_D \stackrel{\text{def}}{=} \mathbf{R}^1 f_^F \{(\tau_{X^{\log}/S^{\log}})^F(D)\}$. Let $a : S \rightarrow \mathcal{A}$ be a section of $\mathcal{A} \rightarrow S$. Then the \mathbf{P}^1 -bundle with logarithmic connection $(P_a \rightarrow X, \nabla_{P_a})$ constructed as above from a is an (i, j) -link bundle with p -curvature locus D – hence, in particular, crys-stable – and, moreover, any (i, j) -link bundle on X^{\log} with p -curvature locus D is of the form $(P_a \rightarrow X, \nabla_{P_a})$.*

The explicit forms of (i, j) -link bundles given in Propositions 1.4 and 1.5 allow us to prove the following:

Corollary 1.6. *The algebraic stack $\mathcal{R}_{g,r}^{i,j}$ is smooth of dimension $3g-3+r+2(\chi-i)$ (respectively, $2(3g-3+r)$) over \mathbf{F}_p if $i > 0$ (respectively, $i = 0$).*

Proof. We work over an \mathbf{F}_p -scheme S , endowed with a closed subscheme $S_0 \subseteq S$ defined by a square-nilpotent ideal. We suppose that we are given an (i, j) -link bundle $(P_0 \rightarrow X_0, \nabla_{P_0})$ on X_0^{\log} (where $f_0^{\log} : X_0^{\log} \rightarrow S_0^{\log}$ is an r -pointed stable log-curve of genus g), and we ask if we can lift this pair (consisting of an (i, j) -link bundle and a curve) to S . Note first of all that (once the log structure on S^{\log} is fixed) no matter what lifting $f^{\log} : X^{\log} \rightarrow S^{\log}$ of $f_0^{\log} : X_0^{\log} \rightarrow S_0^{\log}$ we choose, the Frobenius conjugate $(f^{\log})^F : (X^{\log})^F \rightarrow S^{\log}$ depends only on f_0^{\log} .

Thus, in the case $j = 0$, we can take any lifting $Q \rightarrow X^F$ of $Q_0 \rightarrow X_0^F$ (in the notation of Proposition 1.4), and pulling back $Q \rightarrow X^F$ via $\Phi_{X/S}$ will give us a crystal $(P; \nabla_P)$ in \mathbf{P}^1 -bundles on $\text{Crys}(X_0^{\log}/S^{\log})$ that is independent of the choice of f^{\log} and lifts the crystal defined by

$(P_0; \nabla_{P_0})$. (Note that such a lifting $Q \rightarrow X^F$ always exists – at least Zariski locally on S – since $X^F \rightarrow S$ has relative dimension one.)

In the case $j > 0$, we can always lift the j -balanced divisor $(D_{\mathcal{P}})_0 \subseteq X_0^F$ to a j -balanced divisor $D_{\mathcal{P}} \subseteq X^F$; then form the torsor $\mathcal{A} \rightarrow S$ of Proposition 1.5, and finally, lift the section $a_0 : S_0 \rightarrow \mathcal{A}$ defined by $(P_0; \nabla_{P_0})$ to a section $a : S \rightarrow \mathcal{A}$ (since $\mathcal{A} \rightarrow S$ is smooth). Thus, we again obtain a crystal $(P; \nabla_P)$ in \mathbf{P}^1 -bundles on $\text{Crys}(X_0^{\log}/S^{\log})$ that is independent of the choice of f^{\log} and lifts the crystal defined by $(P_0; \nabla_{P_0})$. Note that so far there are precisely $3g - 3 + r$ dimensions (i.e., the rank of $\mathbf{R}^1(f_0)_* \text{Ad}(Q_0)$ in the case $j = 0$; the rank of Θ_D plus the degree of D in the case $j > 0$) worth of ambiguity in the choice of the lifting $(P; \nabla_P)$.

The only thing that remains is to check that it is possible to find a lifting $f^{\log} : X^{\log} \rightarrow S^{\log}$ such that when $(P_0; \nabla_{P_0})$ is evaluated on X^{\log} , the resulting bundle is crys-stable of level i . If $i = 0$, there is no problem (that is, any f^{\log} will do). When $i > 0$, the existence of such an f^{\log} follows from Proposition 1.7 of Chapter I (cf. the argument in Lemma 3.8 of Chapter I). The computation of the dimension of $\mathcal{R}_{g,r}^{i,j}$ then follows immediately as in Lemma 3.8 of Chapter I. \circ

Remark. Note that when $i = \chi$, the stack $\mathcal{R}_{g,r}^{i,j}$ naturally embeds as a (locally closed) algebraic substack of $\overline{\mathcal{N}}_{g,r}$. Namely, it is the locally closed substack of $\overline{\mathcal{N}}_{g,r}$ consisting of indigenous bundles whose p -curvatures are of a certain specified type. Thus, we obtain the result that this substack is smooth of dimension $3g - 3 + r$.

§1.3. The Stack Associated to a VF-Pattern

Let $\Pi : \mathbf{Z} \rightarrow \mathfrak{L}\mathfrak{e}\mathfrak{v}$ be a VF-pattern of period ϖ . Let us consider the product (over \mathbf{F}_p)

$$\mathcal{W} \stackrel{\text{def}}{=} \prod_{i=0}^{\varpi-1} \mathcal{R}_{g,r}^{\Pi(i), \Pi(i-1)}$$

We would like to define a relative functor $\mathcal{X} \rightarrow \mathcal{W}$ over \mathcal{W} as follows: Let S be an \mathbf{F}_p -scheme, and consider an S -valued point α of \mathcal{W} . This amounts to giving the following data:

- (1) A total of ϖ (possibly distinct) r -pointed stable log-curves of genus g , $f_i^{\log} : X_i^{\log} \rightarrow S_i^{\log}$ (for $i = 0, \dots, \varpi - 1$), where the underlying scheme of S_i^{\log} is S . One can think of the log structure on S_i^{\log} as that obtained by pulling back to S via the classifying morphism of f_i the log structure of $\overline{\mathcal{M}}_{g,r}^{\log}$.

- (2) On each X_i^{\log} , a $(\Pi(i), \Pi(i-1))$ -link bundle (P_i, ∇_{P_i}) , which we call the i^{th} link bundle of α .

Then a point β of \mathcal{X} lying over α consists of giving (in addition) the following data: For each dormant i such that $0 \leq i < \varpi$, we are given

- (1) an isomorphism of log-curves $\xi_{X_i}^{\log} : X_i^{\log} \cong X_{i+1}^{\log}$ (where we regard a subscript ϖ as a subscript 0); we shall call these isomorphisms the *curve-link isomorphisms* of β ;
- (2) a horizontal isomorphism ξ_{P_i} (in the category of \mathbf{P}^1 -bundles with connection on X_i^{\log}) of $\Phi_{X_i}^* P_i$ (equipped with the connection for which sections of $\Phi_{X_i}^{-1} P_i$ are horizontal) with $\xi_{X_i}^*(P_{i+1}, \nabla_{P_{i+1}})$; we shall call these isomorphisms the *bundle-link isomorphisms* of β .

Then we have the following result:

Proposition 1.7. *The morphism $\mathcal{X} \rightarrow \mathcal{W}$ is representable by a finite and unramified scheme over \mathcal{W} . Moreover, \mathcal{X} is a smooth stack over \mathbf{F}_p of dimension*

$$\sum_i \{3g - 3 + r + 2(\chi - \Pi(i))\}$$

where the sum is over active i such that $0 \leq i < \varpi$.

Proof. The first sentence follows from Proposition 1.5 of Chapter I, together with the fact that $\overline{\mathcal{M}}_{g,r}$ is an algebraic stack. The proof of the second sentence follows from Corollary 1.6 by tracing through the definitions (i.e., the extra data (1) and (2) in the above definition of \mathcal{X} imply that it suffices to count only the active “ j ’s”). \circ

Next, we would like to define a relative functor

$$\overline{\mathcal{N}}_{g,r}^{\Pi} \rightarrow \mathcal{X}$$

as follows: Let β be an S -valued point of \mathcal{X} as above. Then an S -valued point γ of $\overline{\mathcal{N}}_{g,r}^{\Pi}$ lying over β consists of giving the following additional data: For each active i such that $0 \leq i < \varpi$, an isomorphism of log-curves $\xi_{X_i}^{\log} : X_i^{\log} \cong X_{i+1}^{\log}$ such that the following conditions are satisfied

- (1) The automorphism of X_0^{\log} given by composing all the $\xi_{X_i}^{\log}$ ’s is the identity. This means that we can, without

confusion, regard (by means of the curve- and bundle-link isomorphisms) all the link bundles as bundles on X_0^{\log} .

- (2) Let $D_i \subseteq X_i$ be the Kodaira-Spencer locus of P_i . Then $\xi_{X_i}(D_i)^F \subseteq X_{i+1}^F$ (where the superscript “ F ” denotes pull-back with respect to the Frobenius morphism $\Phi_S : S \rightarrow S$) is the p -curvature locus of $(P_{i+1}, \nabla_{P_{i+1}})$.

Note that by assigning to the above point the log-curve X_0^{\log} , we obtain a natural projection

$$\overline{\mathcal{N}}_{g,r}^{\Pi} \rightarrow \overline{\mathcal{M}}_{g,r}$$

Now we have the following:

Lemma 1.8. *The morphism $\overline{\mathcal{N}}_{g,r}^{\Pi} \rightarrow \mathcal{X}$ is representable by a finite and unramified scheme over \mathcal{W} . Moreover, $\overline{\mathcal{N}}_{g,r}^{\Pi}$ can étale locally be written as the closed subscheme of a smooth scheme of dimension $\dim_{\mathbf{F}_p}(\mathcal{X})$ over \mathbf{F}_p defined by a set of $\dim_{\mathbf{F}_p}(\mathcal{X}) - (3g - 3 + r)$ equations. In particular, $\dim_{\mathbf{F}_p}(\overline{\mathcal{N}}_{g,r}^{\Pi}) \geq 3g - 3 + r$.*

Proof. The first sentence follows from the fact that $\overline{\mathcal{M}}_{g,r}$ is an algebraic stack, plus the fact that the condition (2) is a condition to the effect that two divisors on a curve be equal, hence is satisfied precisely over some closed subscheme of the base. The second sentence follows from the fact that $\overline{\mathcal{N}}_{g,r}^{\Pi} \rightarrow \mathcal{X}$ is unramified, plus counting the number of conditions. \circ

We are now ready to prove the analogue for the stacks $\overline{\mathcal{N}}_{g,r}^{\Pi}$ of Theorem 1.10 of Chapter II: Let $(\overline{\mathcal{N}}_{g,r}^{\Pi})' \subseteq \overline{\mathcal{N}}_{g,r}^{\Pi}$ be the open substack whose k -valued points (where k is an algebraically closed field) are those for which all the Kodaira-Spencer loci $D_i \subseteq X_i$ of all the link bundles involved satisfy the following condition:

(*) The multiplicity in D_i of every $x \in X_i(k)$ is $\leq p - 2$.

(cf. Theorem 1.9 of Chapter II).

Theorem 1.9. *Suppose that the VF-pattern Π contains no zeroes. Then the morphism $(\overline{\mathcal{N}}_{g,r}^{\Pi})' \rightarrow \overline{\mathcal{M}}_{g,r}$ is flat and quasi-finite. Moreover, $(\overline{\mathcal{N}}_{g,r}^{\Pi})'$ is a local complete intersection of dimension $3g - 3 + r$ over \mathbf{F}_p .*

Proof. Since $\overline{\mathcal{M}}_{g,r}$ is regular, it follows from the second sentence of Lemma 1.8 that it suffices to prove that $(\overline{\mathcal{N}}_{g,r}^{\Pi})' \rightarrow \overline{\mathcal{M}}_{g,r}$ is quasi-finite.

Thus, it suffices to consider points γ of $(\overline{\mathcal{N}}_{g,r}^\Pi)'$ over an algebraically closed field k . We consider the curve X_0^{\log} to be fixed, and we would like to show that there are only a finite number of possibilities for γ . Also, we may identify all the X_i^{\log} 's with X_0^{\log} ; we will denote the resulting log-curve by X^{\log} . Observe that the Kodaira-Spencer locus of (P_0, ∇_{P_0}) is the empty set, hence is always completely determined.

Now suppose that the Kodaira-Spencer locus of (P_b, ∇_{P_b}) has been determined (for some $b \in \mathbf{Z}$). Theorem 1.9 of Chapter II implies that there is only a finite number of possibilities for (P_b, ∇_{P_b}) . Then it follows from condition (2) of the definition of $\overline{\mathcal{N}}_{g,r}^\Pi$ that the Kodaira-Spencer locus of $(P_{b-1}, \nabla_{P_{b-1}})$ is completely determined by (P_b, ∇_{P_b}) . Thus, (up to replacing b by $b - 1$) we see that we are back to the situation postulated at the beginning of this paragraph. Continuing in this fashion, we see that all the data of γ are completely determined, up to a finite number of possibilities. This completes the proof of the Theorem. \circ

We now make the following definition:

Definition 1.10. We shall call $\overline{\mathcal{N}}_{g,r}^\Pi$ the *stack associated to the VF-pattern Π of period ϖ* , or, more succinctly, the *VF-stack (of type (Π, ϖ))*.

Note that even though we omitted “ ϖ ” from the notation $\overline{\mathcal{N}}_{g,r}^\Pi$ (to keep the number of indices to a minimum), the stack $\overline{\mathcal{N}}_{g,r}^\Pi$ still, nonetheless, depends on ϖ as well as g , r , and Π .

Remark. The stacks $\overline{\mathcal{N}}_{g,r}^\Pi$ play an important role in this book, comparable to the role of $\overline{\mathcal{N}}_{g,r}$ in [Mzk1]: Namely, it is over (certain open substacks of) these stacks over that the uniformization theory will actually take place. In some sense, (at least in the case that the VF-pattern Π contains no zeroes) all the information contained in $\overline{\mathcal{N}}_{g,r}^\Pi$ is already contained in the stacks $\overline{\mathcal{N}}_{g,r}^{\rho;l}$ of Chapter II. Unfortunately, however, the form in which this information is contained in $\overline{\mathcal{N}}_{g,r}^{\rho;l}$ turns out not to be very useful or natural from the point of view of the theory to be discussed in this book. This fact will grow more and more obvious as the reader progresses through this book. On the other hand, the stacks $\overline{\mathcal{N}}_{g,r}^\Pi$ are very natural from the point of view of this book, and, moreover, possess other interesting properties such as the affineness properties discussed below in §2.

Finally, before proceeding, we would like to introduce a slight variant of the above VF-stack, defined as follows: We begin with the same product (over \mathbf{F}_p) as above:

$$\mathcal{W} \stackrel{\text{def}}{=} \prod_{i=0}^{\varpi-1} \mathcal{R}_{g,r}^{\Pi(i), \Pi(i-1)}$$

We would like to define a relative functor $\overline{\mathcal{N}}_{g,r}^{\Pi,s} \rightarrow \mathcal{W}$ over \mathcal{W} as follows: Let S be an \mathbf{F}_p -scheme, and consider an S -valued point α of \mathcal{W} . This amounts to giving the following data:

- (1) A total of ϖ (possibly distinct) r -pointed stable log-curves of genus g , $f_i^{\log} : X_i^{\log} \rightarrow S_i^{\log}$ (for $i = 0, \dots, \varpi - 1$), where the underlying scheme of S_i^{\log} is S .
- (2) On each X_i^{\log} , a $(\Pi(i), \Pi(i-1))$ -link bundle (P_i, ∇_{P_i}) , which we call the i^{th} link bundle of α .

Then a point β of $\overline{\mathcal{N}}_{g,r}^{\Pi,s}$ lying over α consists of giving (in addition) data (subject to various conditions) as follows:

- (1) For each nonindigenous i such that $0 < i < \varpi$, we are given an isomorphism of log-curves $\xi_{X_i}^{\log} : X_i^{\log} \cong (X_{i+1}^{\log})^F$ (where we regard a subscript ϖ as a subscript 0). For each Π -ind-adjacent (i, j) such that $0 < i, j \leq \varpi$, we are given an isomorphism of log-curves $\xi_{X_i}^{\log} : X_i^{\log} \cong X_j^{\log}$. We shall call these isomorphisms the *curve-link isomorphisms* of β . Note that by means of these isomorphisms, every X_i^{\log} can be identified with some Frobenius conjugate of X_0^{\log} .
- (2) For each dormant i such that $0 < i < \varpi$, a horizontal isomorphism ξ_{P_i} of $\Phi_{X_{i+1}/S}^*(\xi_{X_i})_*(P_i, \nabla_{P_i})$ with $(P_{i+1}, \nabla_{P_{i+1}})$; we shall call these isomorphisms the *bundle-link isomorphisms* of β .
- (3) For each active, nonindigenous i such that $0 \leq i < \varpi$, let $D_i \subseteq X_i$ be the Kodaira-Spencer locus of P_i . Then $(\xi_{X_i})_*(D_i) \subseteq X_{i+1}^F$ is the p -curvature locus of $(P_{i+1}, \nabla_{P_{i+1}})$.

Note that by assigning to the above point the log-curve X_0^{\log} , we obtain a natural projection

$$\overline{\mathcal{N}}_{g,r}^{\Pi,s} \rightarrow \overline{\mathcal{M}}_{g,r}$$

Moreover, it follows as above that $\overline{\mathcal{N}}_{g,r}^{\Pi,s}$ is representable by an algebraic stack of finite type over \mathbf{F}_p .

Definition 1.11. We shall refer to $\overline{\mathcal{N}}_{g,r}^{\Pi,s}$ as the *shifted VF-stack of type (Π, ϖ)* .

§2. Affineness Properties

§2.1. A Trivialization of a Certain Line Bundle on $\overline{\mathcal{N}}_{g,r}^{\Pi}$

Let $S = \overline{\mathcal{N}}_{g,r}^{\Pi}$. Let $f^{\log} : X^{\log} \rightarrow S^{\log}$ be the tautological r -pointed log curve of genus g over S (where the log structure on S is the pull-back of that on $\overline{\mathcal{M}}_{g,r}$). Thus, we have tautological divisors $D_i \subseteq X$, and tautological link bundles (P_i, ∇_{P_i}) on X^{\log} . Let $\mathcal{T}_i = \tau_{X^{\log}/S^{\log}}(D_i)$. Thus, \mathcal{T}_i is a line bundle on X . Suppose that (a, b) is “ Π -adjacent” (cf. §1.1 for a definition of this term). Then by pulling back the p -curvature by Frobenius $(b - a - 1)$ times, we have a locally split morphism

$$(\Phi_X^*)^{b-a} \mathcal{T}_a \hookrightarrow \mathrm{Ad}(P_b)$$

On the other hand, the Hodge filtration on $\mathrm{Ad}(P_b)$ defines a surjection

$$\mathrm{Ad}(P_b) \rightarrow \mathcal{T}_b$$

Composing these two morphisms, we obtain a morphism of line bundles

$$\mathcal{H}_{(a,b)} : (\Phi_X^*)^{b-a} \mathcal{T}_a \rightarrow \mathcal{T}_b$$

which is a sort of generalization of the (square of the) Hasse invariant. We shall refer to $\mathcal{H}_{(a,b)}$ as the (a, b) -square Hasse invariant of the data under consideration. Note that since $\mathcal{H}_{(a,b)}$ is *nonzero generically* on every geometric irreducible component of every fiber of $X \rightarrow S$ (indeed, this is a consequence of the fact that the (P_i, ∇_{P_i}) are all *crys-stable*), it follows that its zero locus $\mathcal{D}_{(a,b)} \subseteq X$ is S -flat. It is not difficult to see (cf. Proposition 2.6 of Chapter II of [Mzk1]) that the divisor $\mathcal{D}_{(a,b)}$ can, in fact, be written in the form $2\mathcal{D}_{(a,b)}^{\mathrm{rd}}$, where $\mathcal{D}_{(a,b)}^{\mathrm{rd}} \subseteq X$ is an S -flat divisor. Moreover, when restricted to $\mathcal{D}_{(a,b)}^{\mathrm{rd}}$, the morphism $(\Phi_X^*)^{b-a} \mathcal{T}_a \hookrightarrow \mathrm{Ad}(P_b)$ maps into the subbundle $(\mathcal{T}_b)^{\vee} \subseteq \mathrm{Ad}(P_b)$ arising from the Hodge section. That is to say, we have an isomorphism

$$\mathcal{K}_{(a,b)} : (\Phi_X^*)^{b-a} \mathcal{T}_a|_{\mathcal{D}_{(a,b)}^{\mathrm{rd}}} \cong (\mathcal{T}_b)^{\vee}|_{\mathcal{D}_{(a,b)}^{\mathrm{rd}}}$$

We would like to do some arithmetic with $\mathcal{H}_{(a,b)}$ and $\mathcal{K}_{(a,b)}$ in order to prove that a certain line bundle on $\overline{\mathcal{N}}_{g,r}^{\Pi}$ is trivial.

To do this, we would like to make use of Chow groups (with \mathbf{Q} -coefficients) on X and S . We refer to [Fulton1] for basic facts on Chow groups on quasi-projective schemes. For the reader who is squeamish about working with Chow groups of stacks, we note that for the two stacks in question (X and S), the associated coarse moduli spaces are quasi-projective schemes. Thus, since we are always working with \mathbf{Q} -coefficients, questions about the Chow groups of the stacks in question can always be easily reduced to questions concerning the Chow group of quasi-projective schemes. We shall denote the Chow group (with \mathbf{Q} -coefficients) generated by primes of codimension d on X by $\mathrm{CH}^d(X)$. Let $\tau_i = c_1(\mathcal{T}_i) \in \mathrm{CH}^1(X)$ be the first Chern class of \mathcal{T}_i . Let $\eta_i \in \mathrm{CH}^1(S)$ be the class given by $f_*(\tau_i^2)$, where $f_* : \mathrm{CH}^2(X) \rightarrow \mathrm{CH}^1(S)$ is the push-forward morphism (which is well-defined since $X \rightarrow S$ is proper). The class $[\mathcal{D}_{(a,b)}^{\mathrm{rd}}] \in \mathrm{CH}^1(X)$ defined by the divisor $\mathcal{D}_{(a,b)}^{\mathrm{rd}} \subseteq X$ is thus given by $\frac{1}{2}(\tau_b - p^{b-a}\tau_a)$. On the other hand, the existence of $\mathcal{K}_{(a,b)}$ implies that if we multiply $\mathcal{D}_{(a,b)}^{\mathrm{rd}}$ by $\tau_b + p^{b-a}\tau_a$, we get zero. That is to say,

$$\tau_b^2 - p^{2(b-a)}\tau_a^2 = 0$$

in $\mathrm{CH}^2(X)$. Applying f_* to this equation, we thus, obtain that

$$\eta_b - p^{2(b-a)}\eta_a = 0$$

in $\mathrm{CH}^1(S)$.

So far in the discussion, the Π -adjacent pair (a, b) has been kept fixed. However, if we allow the pair (a, b) to vary over all Π -adjacent pairs in one period of Π , and we add up the resulting equations (with appropriate coefficients – which will necessarily be powers of p), we obtain that for every active integer a , η_a is annihilated by some nonzero integer. (In fact, it is not difficult to see that this integer is a p -adic unit. However, we shall not use this fact.) Since $\mathrm{CH}^1(S)$ is a \mathbf{Q} -vector space, it thus follows that $\eta_a = 0$, for all such a . In summary,

Theorem 2.1. *For each Π -active $i \in \mathbf{Z}$, the class $f_*c_1(\omega_{X/S}^{\log}(-D_i))^2 \in \mathrm{CH}^1(\overline{\mathcal{N}}_{g,r}^{\Pi})$ is zero. In particular, the class $f_*c_1(\omega_{X/S}^{\log})^2 \in \mathrm{CH}^1(\overline{\mathcal{N}}_{g,r}^{\Pi})$ is zero.*

This is of particular interest since the class $\eta_0 = f_*c_1(\omega_{X/S}^{\log})^2$ is the pull-back of a class defined on $\overline{\mathcal{M}}_{g,r}$.

Remark. Note that this class is a rational multiple of the class (computed in Theorem 3.4 of Chapter I of [Mzk1]) associated to the Schwarz torsor. Thus, it is entirely appropriate that it should be zero over $\overline{\mathcal{N}}_{g,r}^{\Pi}$, where the Schwarz torsor admits a trivialization. On the

other hand, what we have proven here is stronger than the statement that the Schwarz torsor is trivial over $\overline{\mathcal{M}}_{g,r}^{\Pi}$: the difference is precisely the difference between saying that a line bundle is trivial and saying that the line bundle admits a connection.

§2.2. Some Ampleness Results

Let $S = \overline{\mathcal{M}}_{g,r}$. Let $f^{\log} : X^{\log} \rightarrow S^{\log}$ be the universal r -pointed log-curve of genus g . Let $\eta_{g,r} \stackrel{\text{def}}{=} f_* c_1(\omega_{X/S}^{\log})^2 \in \text{CH}^1(S)$. Note that the class $\eta_{g,r}$ can always be expressed as a rational multiple of the first Chern class of some line bundle on S . Moreover, there always exists a (nonzero) tensor power of any line bundle on S that descends to the coarse moduli space associated to $\overline{\mathcal{M}}_{g,r}$. Thus, with these remarks made, it makes sense to say that the class $\eta_{g,r}$ is *ample*: that is, we mean that it is a positive rational multiple of the first Chern class of the pull-back to S of an ample line bundle on the coarse moduli space. Now we have the following result:

Proposition 2.2. *The class $\eta_{g,r}$ is ample.*

Proof. The plan is to use induction on $3g - 3 + r$ (the dimension of $\overline{\mathcal{M}}_{g,r}$), and apply Seshadri's criterion for ampleness. For any clutching morphism

$$\prod_i \overline{\mathcal{M}}_{g_i, r_i} \rightarrow \overline{\mathcal{M}}_{g,r}$$

it is easy to see that the restriction of $\eta_{g,r}$ to $\mathcal{P} \stackrel{\text{def}}{=} \prod_i \overline{\mathcal{M}}_{g_i, r_i}$ is given by

$$\sum_i (\eta_{g_i, r_i}|_{\mathcal{P}})$$

Thus, the induction hypothesis implies that the restriction of $\eta_{g,r}$ to the boundary of $\overline{\mathcal{M}}_{g,r}$ is ample. We shall henceforth assume that $3g - 3 + r \geq 1$.

It turns out that it is easier to derive the Proposition from well-known results in the case $r = 0$. Thus, in order to reduce to that case, we make use of the stacks of *log admissible coverings* of [Mzk2]. Let $\mathcal{H}' \rightarrow \overline{\mathcal{M}}_{g,r}$ be a connected component of the stack of (nontrivial) log admissible coverings $Y^{\log} \rightarrow X^{\log}|_{\mathcal{H}'}$ of degree 2 which are étale at all but an even number of marked points. (Here, we denote by $X^{\log} \rightarrow \overline{\mathcal{M}}_{g,r}^{\log}$ the tautological log-curve over $\overline{\mathcal{M}}_{g,r}^{\log}$.) Note that Y^{\log} always has an even number of marked points. Let $\mathcal{H} \rightarrow \mathcal{H}'$ be the stack of log admissible

coverings $Z^{\log} \rightarrow Y^{\log}|_{\mathcal{H}}$ of degree two that are ramified at *all* the marked points. Thus, $\mathcal{H} \rightarrow \overline{\mathcal{M}}_{g,r}$ is finite, flat, and surjective. Moreover, over \mathcal{H} , we have a log admissible covering $Z^{\log} \rightarrow X^{\log}|_{\mathcal{H}}$. Suppose that Z^{\log} is an s -pointed stable log-curve of genus q . Thus, $2q - 2 + s = 4(2g - 2 + r)$. Note that q is always ≥ 2 . Let $Z' \rightarrow \mathcal{H}$ be the stable curve of genus q obtained by forgetting the marked points. Then we *claim* that the natural morphism $\mathcal{H} \rightarrow \overline{\mathcal{M}}_q$ (that assigns to the covering $Z^{\log} \rightarrow X^{\log}|_{\mathcal{H}}$ the curve Z') is *quasi-finite* on the open substack of \mathcal{H} given by the inverse image of $\mathcal{M}_{g,r} \subseteq \overline{\mathcal{M}}_{g,r}$ (via $\mathcal{H} \rightarrow \overline{\mathcal{M}}_{g,r}$). Indeed, over $\mathcal{M}_{g,r}$, X is smooth; hence so is $Z \rightarrow \mathcal{H}$. Thus, over $\mathcal{M}_{g,r}$, Z and Z' coincide, and the covering $Z \rightarrow X|_{\mathcal{H}}$ is determined by a group of automorphisms of Z . On the other hand, since $q \geq 2$, the automorphism group of Z is finite. This proves the claim.

Let us denote by means of a superscript “c” the “coarse moduli space” associated to an algebraic stack (such as \mathcal{H} or $\overline{\mathcal{M}}_{g,r}$). (We refer to [FC], Chapter 1, Theorem 4.10, for more details on the notion of the “associated coarse moduli space.”) We observe in passing that although, in general, the associated coarse moduli space of an algebraic stack can only be shown to be an algebraic space (i.e., it is not necessarily a scheme), in the case of $\overline{\mathcal{M}}_{g,r}$ (and hence also of \mathcal{H} , which is finite over $\overline{\mathcal{M}}_{g,r}$), it is well-known ([Mumf1], Corollary 5.18) that $\overline{\mathcal{M}}_{g,r}^c$ is a scheme. Next, note that since \mathcal{H} and $\overline{\mathcal{M}}_q$ are of finite type over \mathbf{F}_p , it is a consequence of the claim discussed in the preceding paragraph that there exists an integer N such that the following condition is satisfied:

(*) Suppose that $C_{\mathcal{H}} \subseteq \mathcal{H}^c$ is a *curve* (i.e., an irreducible subvariety of dimension one), and denote by $C_q \subseteq \overline{\mathcal{M}}_q^c$ its image in $\overline{\mathcal{M}}_q^c$. Suppose further that the image of $C_{\mathcal{H}}$ in $\overline{\mathcal{M}}_{g,r}^c$ has nonempty intersection with $\mathcal{M}_{g,r}^c \subseteq \overline{\mathcal{M}}_{g,r}^c$. Then the generic degree of $C_{\mathcal{H}}$ over C_q is $\leq N$.

This integer N will play an important role in the remainder of the proof.

Next, we explain how we will use this integer N . If V is a proper variety over \mathbf{F}_p , and $C \subseteq V$ is a curve, then we denote by

$$\mu_V(C) \stackrel{\text{def}}{=} \sup_{P \in C} \{\mu_P(C)\}$$

Here, the “Sup” is over all *closed* points P in C , and $\mu_P(C)$ is the multiplicity of the maximal ideal \mathfrak{m}_P of the local ring $\mathcal{O}_{C,P}$ of C at P . We recall that this *multiplicity* may be defined as the limit

$$\lim \left(\frac{\text{length}(\mathcal{O}_{C,P}/\mathfrak{m}_P^n)}{n} \right)$$

as $n \rightarrow \infty$. Thus, in the context of $(*)$, I *claim* that

$$N \cdot \mu_{\overline{\mathcal{M}}_q^c}(C_q) \geq \mu_{\mathcal{H}^c}(C_{\mathcal{H}})$$

Indeed, if $P_{\mathcal{H}}$ is a closed point of $C_{\mathcal{H}}$ that maps to a closed point P_q of C_q , then the associated morphism of complete local rings $A \stackrel{\text{def}}{=} \widehat{\mathcal{O}}_{C_q, P_q} \rightarrow B \stackrel{\text{def}}{=} \widehat{\mathcal{O}}_{C_{\mathcal{H}}, P_{\mathcal{H}}}$ is finite. Moreover, the definition of N implies that there exists a morphism of A -modules $\phi: A^N \rightarrow B$ whose cokernel is of finite length over A . Thus, it follows that (for all $n \geq 0$) $\phi(\mathfrak{m}_A^n \cdot A^N) \subseteq \mathfrak{m}_B^n$. In particular,

$$\text{length}(B/\mathfrak{m}_B^n) \leq N \cdot \text{length}(A/\mathfrak{m}_A^n) + \text{const.}$$

where “const.” is a constant independent of n . Letting $n \rightarrow \infty$ thus completes the proof of the claim.

Next, we propose to show that the pull-back of $\eta_{g,r}$ to \mathcal{H} is ample. Since $\mathcal{H} \rightarrow \overline{\mathcal{M}}_{g,r}$ is finite, flat, and surjective, this will complete the proof of the Proposition. Thus, let $C \subseteq \mathcal{H}^c$ be a curve. We would like to *apply Seshadri’s criterion for ampleness*: Thus, we would like to find a constant $\epsilon > 0$ (which is independent of C), but which has the property that for every C as above,

$$\deg(\eta_{g,r}|_C) > \epsilon \cdot \mu_{\mathcal{H}}(C)$$

(see, e.g., [Harts1], p. 37, or [Knud], p. 209, for more details on Seshadri’s criterion). Now on the one hand, if the image of C in $\overline{\mathcal{M}}_{g,r}^c$ lies inside the boundary of $\overline{\mathcal{M}}_{g,r}$, then the existence of such an ϵ follows from the induction hypothesis (i.e., the fact that $\eta_{g,r}$ is ample when restricted to the boundary of $\overline{\mathcal{M}}_{g,r}$). Thus, we may assume that the curve $X_C \rightarrow C$ (where the subscript “ C ” denotes restriction to C) is generically smooth.

Now let us consider the pulled back admissible covering $Z_C \rightarrow X_C$ over C . Let $M_X \subseteq X_C$ (respectively, $M_Z \subseteq Z_C$) denote the divisor of marked points. Let $E \subseteq Z_C$ denote the effective divisor such that $\omega_{Z'_C/C}|_Z(E) = \omega_{Z_C/C}$. Thus, E maps to a finite set of points under the contraction morphism $Z \rightarrow Z'$. Then we compute:

$$\begin{aligned} 4^2 c_1(\omega_{X_C/C}(M_X))^2 &= c_1(\omega_{Z_C/C}(M_Z))^2 \\ &= c_1(\omega_{Z_C/C}(M_Z)) \cdot (c_1(\omega_{Z'_C/C}) + [M_Z] + [E]) \\ &= c_1(\omega_{Z_C/C}(M_Z)) \cdot c_1(\omega_{Z'_C/C}) + c_1(\omega_{Z_C/C}(M_Z)) \cdot ([M_Z] + [E]) \\ &= (c_1(\omega_{Z'_C/C}) + [M_Z] + [E]) \cdot c_1(\omega_{Z'_C/C}) + c_1(\omega_{Z_C/C}(M_Z)) \cdot [E] \\ &\geq (c_1(\omega_{Z'_C/C}) + [M_Z] + [E]) \cdot c_1(\omega_{Z'_C/C}) \\ &= (c_1(\omega_{Z'_C/C}) + [M_Z]) \cdot c_1(\omega_{Z'_C/C}) \\ &\geq c_1(\omega_{Z'_C/C})^2 \end{aligned}$$

where the first inequality follows from the fact that $\omega_{Z_C/C}(M_Z)$ is ample on the fibers of $Z_C \rightarrow C$, and the last inequality follows from the numerical positivity of $\omega_{Z'_C/C}$ ([Szp], p. 56, Theorem 2').

Now let us denote by $\lambda_{Z'}$ the determinant of the push-forward of $\omega_{Z'_C/C}$ by $Z'_C \rightarrow C$. Let $\delta_{Z'}$ be the pull-back to C of the divisor at infinity of $\overline{\mathcal{M}}_q$. Then “Noether’s formula” (a special case of Grothendieck-Riemann-Roch – see, e.g., [Szp], p. 48, Lemme 1) implies that

$$c_1(\omega_{Z'_C/C})^2 = 12 \cdot \deg(\lambda_{Z'}) - \deg(\delta_{Z'})$$

On the other hand, [Mumf1], Corollary 5.18, implies that there exists a universal constant $\epsilon_1 > 0$ such that

$$(12 - \epsilon_1) \cdot \deg(\lambda_{Z'}) - \deg(\delta_{Z'}) \geq 0$$

Thus, we obtain that

$$c_1(\omega_{Z'_C/C})^2 \geq \epsilon_1 \cdot \deg(\lambda_{Z'})$$

Next, we observe that there exist universal constants $\epsilon_2, \epsilon_3 > 0$ such that

$$\deg(\lambda_{Z'} - \epsilon_2 \cdot \delta_{Z'}) \geq \epsilon_3 \cdot \mu_{\overline{\mathcal{M}}_q^c}(C_q)$$

where C_q is the image of C in $\overline{\mathcal{M}}_q^c$. Indeed, this follows from applying the necessariness part of Seshadri’s criterion to the line bundle on $\overline{\mathcal{M}}_q$ corresponding to $\lambda_{Z'} - \epsilon_2 \cdot \delta_{Z'}$ (which is known to be ample for some $\epsilon_2 > 0$ by [Mumf1], Corollary 5.18). Since $\deg(\delta_{Z'}) \geq 0$, we thus obtain:

$$\begin{aligned} 16 \deg(\eta_{g,r}|_C) &= 4^2 c_1(\omega_{X_C/C}(M_X))^2 \\ &\geq c_1(\omega_{Z'_C/C})^2 \\ &\geq \epsilon_1 \cdot \epsilon_3 \cdot \mu_{\overline{\mathcal{M}}_q^c}(C_q) \\ &\geq \frac{1}{N} \epsilon_1 \cdot \epsilon_3 \cdot \mu_{\mathcal{H}^c}(C) \end{aligned}$$

where the last inequality follows from the above discussion concerning the integer N . This completes the verification of Seshadri’s criterion for $\eta_{g,r}|_{\mathcal{H}}$, and hence of the Proposition. \circ

Before proceeding, we have another important ampleness result to prove. Let $\Theta \rightarrow \overline{\mathcal{M}}_{g,r}$ be the torsor over $\Phi_{\overline{\mathcal{M}}_{g,r}}^* \Theta_{\overline{\mathcal{M}}_{g,r}^{\log}/\mathbf{F}_p}$ of Proposition

1.5 (i.e., the torsor denoted “ $\mathcal{A} \rightarrow S$ ” in Proposition 1.5) in the case “ $D = \emptyset$.” Next, let us write the affine torsor $\Theta \rightarrow \overline{\mathcal{M}}_{g,r}$ as $\text{Spec}(\mathcal{R})$, where \mathcal{R} is a sheaf of quasi-coherent algebras on $\overline{\mathcal{M}}_{g,r}$. Thus, \mathcal{R} is, étale locally on $\overline{\mathcal{M}}_{g,r}$, isomorphic to the sheaf of rings of polynomials in $3g - 3 + r$ variables with coefficients in $\mathcal{O}_{\overline{\mathcal{M}}_{g,r}}$. Let $\text{Lin}(\mathcal{R}) \subseteq \mathcal{R}$ denote the subsheaf of sections of degree ≤ 1 . Thus, $\text{Lin}(\mathcal{R})$ is a vector bundle of rank $3g - 2 + r$ on $\overline{\mathcal{M}}_{g,r}$. Let

$$\mathbf{P}_\Theta \rightarrow \overline{\mathcal{M}}_{g,r}$$

be the projective bundle given by $\mathbf{P}(\text{Lin}(\mathcal{R}))$. Thus, we have a natural open immersion $\Theta \hookrightarrow \mathbf{P}_\Theta$ over $\overline{\mathcal{M}}_{g,r}$, and \mathbf{P}_Θ is proper over \mathbf{F}_p . Note that the complement of Θ in \mathbf{P}_Θ is a divisor $D_\Theta \subseteq \mathbf{P}_\Theta$ which has the structure of a projective bundle over $\overline{\mathcal{M}}_{g,r}$. In fact, we can even say what it is explicitly, namely, $\mathbf{P}(\Phi_{\overline{\mathcal{M}}_{g,r}}^* \Omega_{\overline{\mathcal{M}}_{g,r}/\mathbf{F}_p}^{\log})$. Now we have the following result:

Lemma 2.3. *The divisor $D_\Theta \subseteq \mathbf{P}_\Theta$ is numerically effective.*

Proof. By definition ([Harts1], p. 34), it suffices to show that for every nonconstant morphism $\phi : C \rightarrow \mathbf{P}_\Theta$ (where C is a smooth, irreducible, proper curve over \mathbf{F}_p), we have $\deg(\phi^* \mathcal{O}_{\mathbf{P}_\Theta}(D_\Theta)) \geq 0$. As in Proposition 2.2, we use induction on $3g - 3 + r$. If the image of C is not contained in D_Θ , then the desired inequality is obvious. Thus, it suffices to consider the case when ϕ maps C into D_Θ . Note that the restriction of $\mathcal{O}_{\mathbf{P}_\Theta}(D_\Theta)$ to $D_\Theta = \mathbf{P}(\Phi_{\overline{\mathcal{M}}_{g,r}}^* \Omega_{\overline{\mathcal{M}}_{g,r}/\mathbf{F}_p}^{\log})$ is simply the $\mathcal{O}(1)$ -bundle on this projective bundle. Thus, it suffices to show that $\Omega_{\overline{\mathcal{M}}_{g,r}/\mathbf{F}_p}^{\log}$ is a *semi-positive vector bundle* on $\overline{\mathcal{M}}_{g,r}$. Since this vector bundle $\Omega_{\overline{\mathcal{M}}_{g,r}/\mathbf{F}_p}^{\log}$ is functorially well-behaved with respect to restriction to the boundary of $\overline{\mathcal{M}}_{g,r}$, it follows from the induction hypothesis on $3g - 3 + r$ that we may assume that the image of C in $\overline{\mathcal{M}}_{g,r}$ is not contained in the boundary of $\overline{\mathcal{M}}_{g,r}$.

Unfortunately, just as was the case in Proposition 2.2, it is easier to derive this result from well-known results in the case when $r = 0$. Thus, we again resort to stacks of admissible coverings to reduce to the case $r = 0$. Thus, let $\mathcal{H} \rightarrow \overline{\mathcal{M}}_{g,r}$ be as in the proof of Proposition 2.2. Over \mathcal{H} , we have two log admissible coverings of degree two: $Z^{\log} \rightarrow Y^{\log}|_{\mathcal{H}}$; $Y^{\log}|_{\mathcal{H}} \rightarrow X^{\log}|_{\mathcal{H}}$. Pulling everything back to C , we have curves $\alpha : X_C \rightarrow C$; $\beta : Y_C \rightarrow C$; and $\gamma : Z_C \rightarrow C$, together with double coverings $\delta : Z_C \rightarrow Y_C$ and $\epsilon : Y_C \rightarrow X_C$. Let M_X , M_Y , and M_Z be the divisors of marked points in X_C , Y_C , and Z_C . Let $\zeta : Z'_C \rightarrow C$ be the curve obtained by forgetting the marked points of Z_C and contracting. Let $\theta : Z_C \rightarrow Z'_C$ be the contraction morphism. Note that we have a pull-back morphism

$$\delta^* : \beta_* \omega_{Y_C/C}^2(M_Y) \rightarrow \gamma_* \omega_{Z_C/C}^2$$

together with a trace morphism

$$\delta_* : \gamma_* \omega_{Z_C/C}^2 \rightarrow \beta_* \omega_{Y_C/C}^2(M_Y)$$

such that the composite $\delta_* \circ \delta^*$ is multiplication by 2. Thus, in particular, δ_* is surjective. Also, we have a pull-back morphism

$$\theta^* : \zeta_* \omega_{Z'_C/C}^2 \rightarrow \gamma_* \omega_{Z_C/C}^2$$

which is generically an isomorphism on C . Thus, by composing δ_* with θ^* , we obtain a morphism of vector bundles on C :

$$\zeta_* \omega_{Z'_C/C}^2 \rightarrow \beta_* \omega_{Y_C/C}^2(M_Y)$$

which is generically surjective. On the other hand, by [Szp], p. 35 (the first “Théorème” – where we take “ m ” to be 1, and “ A ” to be an arbitrary line bundle of positive degree) and p. 56 (Théorème 2’), the vector bundle $\zeta_* \omega_{Z'_C/C}^2$ is *semi-positive*. It thus follows that $\beta_* \omega_{Y_C/C}^2(M_Y)$ is also semi-positive.

Ultimately, however, we wish to prove that $\alpha_* \omega_{X_C/C}^2(M_X)$ is semi-positive. To this end, we observe that just as above, we have a pull-back morphism

$$\epsilon^* : \alpha_* \omega_{X_C/C}^2(M_X) \rightarrow \beta_* \omega_{Y_C/C}^2(M_Y)$$

together with a trace morphism

$$\epsilon_* : \beta_* \omega_{Y_C/C}^2(M_Y) \rightarrow \alpha_* \omega_{X_C/C}^2(M_X)$$

such that the composite $\epsilon_* \circ \epsilon^*$ is multiplication by 2. In particular, ϵ_* is a surjective morphism of vector bundles. Thus, the semi-positivity of $\beta_* \omega_{Y_C/C}^2(M_Y)$ implies that of $\alpha_* \omega_{X_C/C}^2(M_X)$. This completes the proof of the Lemma. \circ

Now we are ready to state and prove the main result of this subsection.

Theorem 2.4. *Let a and b be positive rational numbers. Then the class $a[D_\theta] + b(\eta_{g,r}|_{\mathbf{P}_\Theta}) \in \mathrm{CH}^1(\mathbf{P}_\Theta)$ is ample.*

Proof. The fact that this Theorem is a consequence of Proposition 2.2 and Lemma 2.3 is well-known, but for the convenience of the reader,

we give a complete proof. Let $\gamma \stackrel{\text{def}}{=} a[D_\Theta] + b(\eta_{g,r}|_{\mathbf{P}_\Theta})$. We apply *Nakai's criterion for ampleness* ([Harts1], p. 30). Thus, let $Y \subseteq \mathbf{P}_\Theta$ be a closed irreducible substack of dimension d . It suffices to show that $\gamma^d|_Y > 0$. Let $Z \subseteq \overline{\mathcal{M}}_{g,r}$ be the image of Y in $\overline{\mathcal{M}}_{g,r}$. Let s be the dimension of Z . Then $\eta_{g,r}^i|_Y = 0$ for $i > s$. On the other hand, since $\eta_{g,r}$ is ample on $\overline{\mathcal{M}}_{g,r}$, it can (up to a positive rational multiple) be represented by an effective divisor $E \subseteq \overline{\mathcal{M}}_{g,r}$. Thus, for $i \leq s$, $\eta_{g,r}^i|_Z$ can be represented by a sum (with positive rational coefficients) of generically positioned varieties $W \subseteq Z$ of dimension $s - i$. Let V be the inverse image of W under the morphism $Y \rightarrow Z$. Thus, the dimension of V is $d - i$; hence, $([D_\Theta]^{d-i} \cdot \eta_{g,r}^i)|_Y$ can be written as a sum (with positive rational coefficients) of terms of the form $[D_\Theta]^{d-i}|_V$. But by *Kleiman's theorem* ([Harts1], p. 34) and Lemma 2.3, it follows that $[D_\Theta]^{d-i}|_V \geq 0$. This at least shows that $\gamma^d|_Y \geq 0$. To show that it is, in fact, positive, it suffices to exhibit one term in the above sum which is positive. To this end, let us consider the term where $i = s$. Then $\eta_{g,r}^s|_Z$ can be represented by a nonempty set of points m in $\overline{\mathcal{M}}_{g,r}$. Moreover, V splits up as a union of varieties V_m contained in the fiber of $\mathbf{P}_\Theta \rightarrow \overline{\mathcal{M}}_{g,r}$ over m . That is to say, V_m is a variety in projective space, and $[D_\Theta]^{d-s}|_{V_m}$ is its degree, hence is > 0 . This completes the proof. \circ

§2.3. Affine Stacks

Let A be a noetherian ring. Let S be a separated algebraic stack of finite type over A (see [FC], Chapter 1, §4, for an introductory treatment of such objects). Then, by [FC], Chapter 1, Theorem 4.10, we can always construct an algebraic space S^c , together with a natural morphism $S \rightarrow S^c$, which serves as a “coarse moduli space” for S . See *loc. cit.* for more details. Let \mathcal{L} be a line bundle on S .

Definition 2.5. We shall call S an *affine* (respectively, *quasi-affine*) stack if S^c is affine (respectively, quasi-affine). If some tensor power of \mathcal{L} defines an embedding of S^c into \mathbf{P}_A^N (N -dimensional projective space over A), then we shall say that \mathcal{L} is *ample on S* .

Proposition 2.6. Suppose that S is an affine stack. Let $N \geq 1$ be a natural number such that the order of the automorphism group of every geometric point of S divides N . Let \mathcal{F} be a coherent sheaf on S . Then multiplication by N annihilates $H^i(S, \mathcal{F})$ for all $i > 0$.

Proof. It suffices to compute the higher direct image sheaves $\mathbf{R}^j f_* \mathcal{F}$ for the morphism $f : S \rightarrow S^c$. But étale locally on S^c , f is of the form $\text{Spec}(R)/G \rightarrow \text{Spec}(R^G)$, where G is an automorphism group of a geometric point of S ; R is a ring with G -action; $\text{Spec}(R)/G$ is the quotient of $\text{Spec}(R)$ by G in the sense of stacks (cf., e.g., the discussion of

[Mzk2], §1.2); and R^G is the ring of G -invariants. Moreover, $\mathcal{F}|_{\mathrm{Spec}(R)}$ corresponds to some finite R -module M with G -action, and the higher direct image sheaves in question correspond to the group cohomology modules $H^j(G, M)$. Since these modules are annihilated by N , the Proposition follows immediately from the Leray-Serre spectral sequence for $f : S \rightarrow S^c$ and the fact that S^c is an affine scheme. \circ

Proposition 2.7. *Suppose that T is also a separated algebraic stack of finite type over A . Let $\phi : T \hookrightarrow S$ be an open immersion of stacks. Let us assume that S is proper over A , and that \mathcal{L} is ample on S . Moreover, let us assume that we have a section $s \in \Gamma(S, \mathcal{L})$ which is nonvanishing exactly on T . Then T is an affine stack.*

Proof. Note that $\phi : T \rightarrow S$ induces a natural morphism $\phi^c : T^c \rightarrow S^c$ on the associated coarse moduli spaces which is again an open immersion. Moreover, by raising \mathcal{L} and s to a high power, we can assume that they descend to S^c . This reduces the Proposition to the case where S and T are algebraic spaces. Let us assume for the rest of the proof that S and T are algebraic spaces. Then since \mathcal{L} is ample on S , we see that S is, in fact, a projective scheme, and so T is a quasi-projective scheme. Thus, we are reduced to the case when S and T are schemes. In this case, the Proposition is well-known. \circ

Proposition 2.8. *Suppose that S is a separated algebraic stack over A , and that \mathcal{O}_S is ample on S . Then S is quasi-affine.*

Proof. One reduces immediately to the case where S is an algebraic space, in which case the Proposition is well-known. \circ

Remark. Many other criteria for ampleness (or quasi-ampleness) of schemes (or algebraic spaces) can easily be seen to apply for ampleness of stacks as well by mimicking the proofs of Propositions 2.7 and 2.8.

§2.4. Absolute Affineness

We are now ready to prove the two main results of this §. We begin with a lemma:

Lemma 2.9. *Let Π be a VF-pattern of period ϖ . The natural morphism $\overline{\mathcal{N}}_{g,r}^\Pi \rightarrow \overline{\mathcal{M}}_{g,r}$ is quasi-affine.*

Proof. We begin with an r -pointed stable log-curve of genus g : $X^{\log} \rightarrow S^{\log}$. We examine successively what it means to give various link bundles on X^{\log} as in the definition of $\overline{\mathcal{N}}_{g,r}^\Pi$. The first link bundle is, in the

terminology of [Mzk1], Chapter II, an “FL-bundle” (or, alternatively, in the language of the present book, “a bundle of the type considered in Proposition 1.5, with $D = \emptyset$ ”), hence defines a section of the affine torsor $\Theta \rightarrow \overline{\mathcal{M}}_{g,r}$ (discussed right before Lemma 2.3). The condition that this FL-bundle (which is always, by Proposition 1.5, crys-stable) be of level $\Pi(1)$ is satisfied on a subscheme of the base scheme in question (by Theorem 3.10 of Chapter I). Let $i > 0$ be such that 0 and i are Π -adjacent. Then by pulling back by Frobenius, one sees that this first link bundle determines all the link bundles numbered 1 through i . The condition that the j^{th} link bundle (where $1 \leq j \leq i$) be of level $\Pi(j)$ is, again by Theorem 3.10 of Chapter I, satisfied on a subscheme of the base scheme in question. Once one knows that this i^{th} FL-bundle is of level $\Pi(i)$, one can consider its Kodaira-Spencer locus, which defines a $\Pi(i)$ -balanced divisor on X^{\log} . Then, as in Proposition 1.5, the $(i+1)^{\text{st}}$ link bundle is determined by a section of a certain affine torsor. Thus, we are back in the same situation as when we considered the first link bundle as an FL-bundle. Continuing in this fashion, one sees that the data for a collection of link bundles on X^{\log} that define a point of $\overline{\mathcal{N}}_{g,r}^{\Pi}$ is given by a quasi-affine scheme $T \rightarrow S$ over S . (In fact, $T \rightarrow S$ factors as a composite of immersions and affine torsors.) \circ

Now we put everything together: By Proposition 2.2, the class $\eta_{g,r}$ on $\overline{\mathcal{M}}_{g,r}$ is ample, hence (by Lemma 2.9) its pull-back to $\overline{\mathcal{N}}_{g,r}^{\Pi}$ is ample. On the other hand, by Theorem 2.1, the restriction of the class $\eta_{g,r}$ to $\overline{\mathcal{N}}_{g,r}^{\Pi}$ is zero. Thus, by Proposition 2.8, it follows that $\overline{\mathcal{N}}_{g,r}^{\Pi}$ is *quasi-affine*.

Now suppose (just for the duration of this paragraph) that (Π, ϖ) is the home VF-pattern. Then $\overline{\mathcal{N}}_{g,r}^{\Pi} = \overline{\mathcal{N}}_{g,r}^{\text{adm}}$, i.e., the stack of curves equipped with a nilpotent, admissible indigenous bundle. Since (in the terminology of [Mzk1], Chapter II) such bundles correspond to “FL-bundles” (i.e., bundles of the type considered in Proposition 1.5, with $D = \emptyset$), we have a natural monomorphism $\overline{\mathcal{N}}_{g,r}^{\Pi} \hookrightarrow \Theta$ (where $\Theta \rightarrow \overline{\mathcal{M}}_{g,r}$ is the torsor discussed just before Lemma 2.3). Moreover, by Proposition 1.5, an FL-bundle is necessarily crys-stable, so it follows from Theorem 3.10 of Chapter I that this natural monomorphism $\overline{\mathcal{N}}_{g,r}^{\Pi} \rightarrow \Theta$ is, in fact, proper, hence a closed immersion. On the other hand, we have a compactification $\Theta \hookrightarrow \mathbf{P}_{\Theta}$. We would like to show that $\overline{\mathcal{N}}_{g,r}^{\Pi}$ is *affine*. To do this, let $H \subseteq \mathbf{P}_{\Theta}$ be the closure of the image of the immersion $\overline{\mathcal{N}}_{g,r}^{\Pi} \hookrightarrow \Theta \hookrightarrow \mathbf{P}_{\Theta}$. Thus, we have an open immersion $\overline{\mathcal{N}}_{g,r}^{\Pi} \hookrightarrow H$, and a closed immersion $H \hookrightarrow \mathbf{P}_{\Theta}^{\Pi}$. Let D_{Θ} be the divisor in \mathbf{P}_{Θ} which is the complement of Θ . Let D_H be the restriction of this divisor to H . Thus, $H - D_H = \overline{\mathcal{N}}_{g,r}^{\Pi}$. By Theorem 2.1, the restriction of the class $\eta_{g,r}$ to $\overline{\mathcal{N}}_{g,r}^{\Pi}$ is trivial. For each positive rational number x , let

\mathcal{L}_x be a line bundle on H whose first Chern class is a positive rational multiple of $\eta_{g,r}|_H + x[D_H]$. The fact that $\eta_{g,r}|_{\overline{\mathcal{N}}_{g,r}^\Pi}$ is trivial means that for any x , some positive power of \mathcal{L}_x admits a section over $\overline{\mathcal{N}}_{g,r}^\Pi$ that is nonzero everywhere. Thus, if we take x to be sufficiently large, it follows that some positive power of \mathcal{L}_x admits a section over H that vanishes on D_H , but is nonzero over $\overline{\mathcal{N}}_{g,r}^\Pi$. On the other hand, by Theorem 2.4, \mathcal{L}_x is ample on H . Thus, we conclude from Proposition 2.7 that $\overline{\mathcal{N}}_{g,r}^\Pi$ is *affine*.

Next suppose (just for the duration of this paragraph) that Π is the pre-home VF-pattern of period $\varpi \geq 1$. Then $\overline{\mathcal{N}}_{g,r}^\Pi$ is clearly finite over a product (over \mathbb{F}_p) of copies of $\overline{\mathcal{N}}_{g,r}^{\text{adm}}$. It thus follows that $\overline{\mathcal{N}}_{g,r}^\Pi$ is *affine*. In other words, we have proven the following result:

Theorem 2.10. *Let g, r be nonnegative integers such that $2g - 2 + r \geq 1$. Let Π be a VF-pattern of period ϖ . Then the stack $\overline{\mathcal{N}}_{g,r}^\Pi$ is quasi-affine. If Π is the pre-home VF-pattern of period $\varpi \geq 1$, then the stack $\overline{\mathcal{N}}_{g,r}^\Pi$ is, in fact, affine.*

Remark. We will use Theorem 2.10 in the following subsection to give a new proof of the connectedness of $\overline{\mathcal{M}}_g$. Because of this application, we would like to observe the following: The careful reader may have noticed that in some portions of the proof of Theorem 2.10 (especially in §2.2), we argued as though it is already known that \mathcal{M}_g is connected. In fact, however, we did not use this connectedness in any *substantive* way in the proof of Theorem 2.10. For instance, we could have begun the proofs of the various ampleness results of §2.2 with the phrase “let J be an arbitrary connected component of $\overline{\mathcal{M}}_g$ ” and then proceeded to show ampleness over an arbitrary J . The reason we chose not to do this, however, is that doing so would only have made the notation more complicated, and the essential structure of the proof less transparent. Nevertheless, the reader can easily check that the use of Theorem 2.10 to prove the connectedness of $\overline{\mathcal{M}}_g$ does not involve any “circular reasoning.”

§2.5. The Connectedness of the Moduli Stack of Curves

In this subsection, we note that the ideas developed in this § allow one to give a new proof (at least for p sufficiently large) of the fact that \mathcal{M}_g is connected. First, we make precise what we mean when we say that p must be sufficiently large. Recall the class $\eta_{g,r} \in \text{CH}^1(\overline{\mathcal{M}}_{g,r})$ dealt with above (Proposition 2.2). When $r = 0$, we shall denote this class by η_g . By Proposition 2.2, this class η_g is ample on $\overline{\mathcal{M}}_g$, hence on every connected component C of $\overline{\mathcal{M}}_g$. Thus, the integer $(\eta_g)^{3g-3}|_C$ is positive. Then our assumption on p is the following:

(*) For every integer q such that $2 \leq q \leq g$, and every connected component C of $\overline{\mathcal{M}}_q$, the rational number $(\eta_q)^{3q-3}|_C$ is $\in \mathbf{Z}_p^\times$. Moreover, for any such q , there exists a positive integer N_q prime to p such that the Chow class $N_q \cdot \eta_q$ is the first Chern class of a line bundle on $(\mathcal{M}_q)_{\mathbf{Z}_p}$. Finally, $p > 5^{2g}$.

It is easy to see that, for a given g , (*) is satisfied for sufficiently large p . For the rest of this subsection, we assume that this hypothesis (*) is in force. Note that the last portion of (*) implies that \mathcal{M}_g admits a finite étale covering (obtained by considering level structures on the 3- or 5-torsion points of the Jacobian) which is a scheme and which is of degree prime to p . This means that although, strictly speaking, \mathcal{M}_g is a *stack*, as long as one's coefficients are p -adic, one may manipulate “degrees” (i.e., cycles of codimension zero) on proper components of \mathcal{M}_g as though \mathcal{M}_g were a scheme. Thus, in the following, we shall, without further notice, make use of such degree arguments on proper components of \mathcal{M}_g , as well as on finite coverings of such proper components.

Now we begin the proof. Let k be the algebraic closure of \mathbf{F}_p . We shall work over k ; thus, (just for the rest of this subsection) \mathcal{M}_g will denote the stack of smooth curves of genus g over k . First note that by induction on g ,

It suffices to prove that *no connected component of \mathcal{M}_g is proper*.

Indeed, it follows from the induction hypothesis on g , plus various simple combinatorial considerations (cf. the proof of Theorem 1.12 of Chapter II) that the boundary $\overline{\mathcal{M}}_g - \mathcal{M}_g$ is connected; thus, if we know that every connected component of \mathcal{M}_g touches the boundary, it follows that $\overline{\mathcal{M}}_g$, hence that \mathcal{M}_g is connected.

Thus, let $J \subseteq \mathcal{M}_g$ be a proper connected component. Recall the finite, flat morphism $\mathcal{N}_g \rightarrow \mathcal{M}_g$ of degree p^{3g-3} of [Mzk1], Chapter II, §2 (or, alternatively, Proposition 1.7 of Chapter II of the present work, specialized to the case where “ l ” is χ ; the radii “ ρ ” are all zero; and “ r ” is equal to zero). (Here \mathcal{N}_g is the stack of smooth curves of genus g over k equipped with a nilpotent indigenous bundle.) We restrict this finite, flat morphism to J :

$$\mathcal{N}_J \stackrel{\text{def}}{=} (\mathcal{N}_g)|_J \rightarrow J$$

Now we want to take $I \subseteq (\mathcal{N}_J)_{\text{red}}$ to be an irreducible component such that the p -curvature of the indigenous bundle represented by the generic point of I “vanishes to the maximal possible degree.” More precisely, if there exists an irreducible component $I \subseteq \mathcal{N}_J$ such that the p -curvature of the indigenous

bundle represented by the generic point of I is identically zero, we take such an irreducible component for our I . If not, then for every irreducible component $I \subseteq \mathcal{N}_J$, the p -curvature of the indigenous bundle represented by the generic point of I vanishes precisely on some divisor of the generic tautological curve over J . In this case, we take for our I an irreducible component such that this divisor (on which the generic p -curvature vanishes) is of maximal degree.

Now we claim that the p -curvature of *every* indigenous bundle parametrized by I vanishes to precisely the same degree as the p -curvature of the generic indigenous bundle parametrized by I . Indeed, note that if there were a pair (parametrized by I) consisting of an indigenous bundle (P, ∇_P) on a curve X for which the p -curvature vanished to a greater degree than for the generic indigenous bundle of I , then by the Remark following Corollary 1.6, this pair $\{(P, \nabla_P), X\}$ would be contained in a smooth subvariety K of \mathcal{N}_J of dimension $3g - 3$ with the property that the p -curvature of every indigenous bundle parametrized by K vanishes to precisely the same degree as that of $\{(P, \nabla_P), X\}$. But the existence of this K would contradict the choice of J . This proves the claim. Note, moreover, before we continue, that it follows (again by the Remark following Corollary 1.6) that I is *smooth* over k .

Next we claim that it can never be the case that the indigenous bundle represented by the generic point of I is *admissible*. Indeed, if that were the case, then by the preceding paragraph, every indigenous bundle parametrized by I would be admissible. But the stack of admissible indigenous bundles is precisely \mathcal{N}_g^Π , where Π is the home VF-pattern. Moreover, it follows from Theorem 2.10 that \mathcal{N}_g^Π is affine (and in particular, quasi-affine), hence so is $\mathcal{N}_g^\Pi|_J$. Since I would then be a k -proper quasi-affine substack of $\mathcal{N}_g^\Pi|_J$, we obtain a contradiction. This proves the claim.

Let us go back to the proof that \mathcal{M}_g is connected. We have seen that the generic indigenous bundle of I is not admissible. By Corollary 2.16 of [Mzk1], Chapter II, it thus follows that \mathcal{N}_J is not reduced at the generic point of I . Thus, it follows that:

The degree of I over J is $< p^{3g-3}$.

Moreover, I is smooth and proper over k . Thus, we can consider the crystalline cohomology of I over $W(k)$ (the Witt vectors with coefficients in k). Let us denote the crystalline cohomology of $I/W(k)$ with coefficients in the structure sheaf by $H_{\text{crys}}^i(I)$. Note that we have a natural filtration $F^j(-)$ on $H_{\text{crys}}^i(I)$ induced by the divided powers of the ideals defining the thickenings of I . Thus, for instance, multiplication by p sends $F^j(-)$ into $F^{j+1}(-)$. Let $\zeta_g \in F^1(H_{\text{crys}}^2(I))$ be the first Chern class of $\eta_g|_I$ in the crystalline cohomology of I . Now we have the following

Lemma 2.11. *The class $\zeta_g \in F^2(H_{\text{crys}}^2(I))$.*

Proof. Let $\Sigma \rightarrow J$ be the *geometric Schwarz torsor* over J (i.e., what was denoted “ $\overline{\mathcal{S}}_{g,r}^\rho \rightarrow \overline{\mathcal{M}}_{g,r}$ ” in Chapter I, §4.3, specialized to the case where the radii “ ρ ” are all zero, and “ r ” is equal to zero). Let θ_g be the first Chern class of $\eta_g|_\Sigma$ in the crystalline cohomology of Σ over $W(k)$. Of course, Σ is not proper, so the reader may feel a bit uncomfortable considering the crystalline cohomology of Σ , since it is not clear that it satisfies the properties that one usually expects of crystalline cohomology such as finite-dimensionality. In fact, however, the only properties of this crystalline cohomology that we will use here are that it exists and is “functorial.”

Since the natural morphism $I \rightarrow J$ factors through Σ , and $\zeta_g = \theta_g|_I$, it suffices to show that θ_g belongs to $F^2(H_{\text{crys}}^2(\Sigma))$. To do this, we can use the natural lifting $\Sigma' \rightarrow J'$ (arising from the Schwarz torsor (cf. Chapter I, §4.3) over $(\mathcal{M}_g)_{\mathbf{Z}_p}$) of $\Sigma \rightarrow J$ to stacks that are flat over \mathbf{Z}_p . But then $\eta_g|_J$ is the restriction of a Chow class η'_g on J' . Moreover, there exists a line bundle \mathcal{L} on J' whose first Chern class is an integral, p -adic unit multiple of η'_g . Thus, $\Sigma' \rightarrow J'$ is the torsor of connections on \mathcal{L} . It thus follows that $\mathcal{L}|_{\Sigma'}$ admits a tautological connection. But this implies that its first Chern class $c_1(\mathcal{L}|_{\Sigma'})$ in $H_{\text{crys}}^2(\Sigma')$ lies in $F^2(-)$. On the other hand, the image of $c_1(\mathcal{L}|_{\Sigma'})$ in $H_{\text{crys}}^2(\Sigma)$ is none other than (a p -adic unit multiple of) θ_g . This completes the proof. \circ

Thus, Lemma 2.11 implies that $\zeta_g^{3g-3} \in F^{6(g-1)}(H_{\text{crys}}^{6(g-1)}(I))$. On the other hand, since I is smooth and connected over k , it follows from Poincaré duality that $H_{\text{crys}}^{6(g-1)}(I) \cong W(k)$, where the filtration induced on $W(k)$ is precisely

$$F^n(W(k)) = (p^{n-(3g-3)} \cdot W(k)) \bigcap W(k)$$

for $0 \leq n \leq 6g - 6$. (Indeed, this statement concerning the filtration follows immediately from Theorem 7.2 of [BO], and the fact that $p > 5^{2g} > 6g - 6$.) It thus follows that ζ_g^{3g-3} is divisible by p^{3g-3} . On the other hand, one knows that ζ_g^{3g-3} is a p -adic integer, equal to precisely $\eta_g^{3g-3}|_J$ times the degree of the finite, flat morphism $I \rightarrow J$. Since by (*), $\eta_g^{3g-3}|_J$ is a p -adic unit, we conclude that the degree of the morphism $I \rightarrow J$ is divisible by p^{3g-3} , hence $\geq p^{3g-3}$. But this contradicts the observation concerning the degree of I over J made above. This completes the proof of the following result:

Theorem 2.12. *Under the assumption (*) on p , the stack $(\mathcal{M}_g)_{\mathbf{F}_p}$ (of smooth curves of genus g over \mathbf{F}_p) is geometrically connected over \mathbf{F}_p .*

Remark. As far as the author knows, there are essentially three other proofs of the connectedness of \mathcal{M}_g : The earliest (to the author's knowledge) rigorous proof involves Hurwitz schemes, and is discussed in [Fulton2]. In this case, the technique is essentially analytic (i.e., over \mathbb{C}) and reduces to proving certain combinatorial results about the action of the braid group. Another proof with deep roots is the proof based on classical Teichmüller theory over the complex numbers (cf. [DM], [Bers]). To the author's knowledge, the only purely algebraic proof is that given by [FM] involving admissible coverings. We believe that the proof given above of Theorem 2.12 is the first proof of this fact that depends on characteristic p methods in an essential way. In a way it is very philosophically satisfying in that it is in accordance with the author's philosophical claim that the uniformization theory discussed in this book constitutes a p -adic version of Bers theory, and in particular, that a certain stack (to be constructed in subsequent Chapters and based on the stacks $\overline{\mathcal{N}}_{g,r}^{\Pi}$) is the p -adic analogue of Bers/Teichmüller space. Thus, one may regard the above proof as the p -adic analogue of the proof of the connectedness of \mathcal{M}_g by means of classical Teichmüller theory.

Chapter IV: Construction of Examples

§0. Introduction

So far, we have discussed *VF-stacks* (cf. Chapter III, Definition 1.10) in great generality. The reader who has studied the theory of [Mzk1] will be familiar with various aspects of the VF-stack associated to the home VF-pattern. We shall refer to this VF-stack as the *home VF-stack*. As the name “home” implies, this VF-stack is rather special and fundamental, and it is the main topic of [Mzk1]. On the other hand, the point of the present book is to extend the theory of [Mzk1] to arbitrary VF-stacks. Unfortunately, at this point, it may not even be clear to the reader that there exist interesting (or indeed any!) VF-stacks other than the home VF-stack. Thus, the purpose of this Chapter is to indulge in a bit of zoology and discuss (and in particular, prove the existence of) several important examples of VF-stacks other than the home VF-stacks. Here we will make essential use of the theories of pseudo-torally crys-stable bundles, n -connections, and mildly spiked bundles discussed in Chapter II to construct explicit examples.

Before proceeding, we make the following remark: It may occur to the reader that a rather extraordinary number of pages in Chapter II and the present Chapter is devoted to a fairly *ad hoc* portion of the theory whose sole purpose is to prove certain very specific existence results (Theorems 2.9 and 3.7). The author’s response would be the following: First of all, as already stated, even though these results require many pages to prove, knowing that the theory developed is not vacuous is nevertheless of great importance in this book. Secondly, the reason that these results require so many pages to prove is not because they are really that difficult, but because the author lacks the appropriate language and technology to express in a few pages the relatively simple and compact ideas necessary to prove them. (A typical case is the rather lengthy proof (probably about 20 pages in all!) of the existence of spiked VF-stacks. If the author could only present

the proof to the reader by drawing and erasing pictures on a blackboard, the proof would probably seem very short and transparent.) Thirdly, and most importantly, even though we chose to eventually prove very specific results such as Theorems 2.9 and 3.7, the thrust of the theory of the present Chapter is to show how one can *reduce the issue of producing new examples to an essentially combinatorial problem*, so that the reader can construct as many examples as he (respectively, she) likes by himself (respectively, herself). It is in this sense that the author hopes that the rather large number of pages effectively devoted to proving Theorems 2.9 and 3.7 can be justified.

§1. Explicit Computation in the Case $g=1$; $r=1$; $p=5$

We begin this Chapter by computing $\overline{\mathcal{N}}_{1,1}^\rho$ explicitly in the case $g = 1$, $r = 1$, $p = 5$. We will not use the results of this § in the rest of this book (in particular, we will not use them in the proofs of the main existence results of this Chapter – Theorems 2.9 and 3.7), but we include this computation nonetheless since it provides an interesting concrete example of the theory. For simplicity, we work over k , the algebraic closure of \mathbf{F}_5 .

§1.1. Irreducible Components of Degree Two

The squares in \mathbf{F}_5 are $\{\pm 1\}$. Thus, $\rho^2 = 0$ or ± 1 . Let \mathcal{L} be the line bundle $\Omega_{\overline{\mathcal{M}}_{1,1}}^{\log}$ on $\overline{\mathcal{M}}_{1,1}$. Let $\mathcal{B} \rightarrow \overline{\mathcal{M}}_{1,1}$ be the geometric vector bundle corresponding to \mathcal{L} , i.e.,

$$\mathcal{B} = \mathrm{Spec}(\oplus_{i \geq 0} \mathcal{L}^{-i})$$

Note that since $H^i(\overline{\mathcal{M}}_{1,1}, \mathcal{L}) = 0$ for $i = 0, 1$ (cf., e.g., [KM], Chapter 8, §8.1.7), it follows that for all ρ , there is a canonical isomorphism $\mathcal{B} \cong \overline{\mathcal{S}}_{1,1}^\rho$. In the following, we shall therefore identify $\overline{\mathcal{S}}_{1,1}^\rho$ with \mathcal{B} . Let $c_4 \in \Gamma(\overline{\mathcal{M}}_{1,1}, \mathcal{L}^2)$ and $c_6 \in \Gamma(\overline{\mathcal{M}}_{1,1}, \mathcal{L}^3)$ be the usual modular forms (as in, say, [KM], Chapter 8). Let $\mathcal{D} \rightarrow \overline{\mathcal{M}}_{1,1}$ be the double covering obtained by taking a square root of c_4 . Since c_4 has only simple zeroes (cf., e.g., [KM], Chapter 8, §8.4.1), it follows that \mathcal{D} is a smooth, proper algebraic stack of dimension one. Let $d_4 \in \Gamma(\mathcal{D}, \mathcal{L}|_{\mathcal{D}})$ be a square root of c_4 . For any $\nu \in k$, we may regard $\nu \cdot d_4$ as a section of $\mathcal{B} \rightarrow \overline{\mathcal{M}}_{1,1}$, hence of $\overline{\mathcal{S}}_{1,1}^\rho \rightarrow \overline{\mathcal{M}}_{1,1}$. The *Verschiebung* (cf. Chapter II, §1.3) of this section will be a section ν_ν of \mathcal{L}^5 over \mathcal{D} . Thus, if the section of $\overline{\mathcal{S}}_{1,1}^\rho \rightarrow \overline{\mathcal{M}}_{1,1}$ corresponding to $\nu \cdot d_4$ maps into $\overline{\mathcal{N}}_{1,1}^\rho \subseteq \overline{\mathcal{S}}_{1,1}^\rho$ at the two points of \mathcal{D} lying over $\infty \in \overline{\mathcal{M}}_{1,1}$, then we conclude that ν_ν is a section of $\mathcal{L}^5(-\infty) \cong \mathcal{L}^{-1}$ (note: this isomorphism follows from the fact that the “discriminant

modular form" Δ has a simple zero at ∞ – cf. [KM], Chapter 8, §8.1.7) over \mathcal{D} , hence will vanish (since \mathcal{L} is *ample*). That is, we have the following:

Lemma 1.1. *The section $\nu \cdot d_4 : \mathcal{D} \rightarrow \overline{\mathcal{S}}_{1,1}^\rho$ maps into $\overline{\mathcal{N}}_{1,1}^\rho$ over all of \mathcal{D} if and only if it does at the two points of \mathcal{D} lying over $\infty \in \overline{\mathcal{M}}_{1,1}$.*

Now let us suppose that there exists an irreducible component $\mathcal{I} \subseteq (\overline{\mathcal{N}}_{1,1}^\rho)_{\text{red}}$ of degree two over $\overline{\mathcal{M}}_{1,1}$. The inclusion $\mathcal{I} \subseteq \overline{\mathcal{N}}_{1,1}^\rho \subseteq \overline{\mathcal{S}}_{1,1}^\rho = \mathcal{B}$ defines a section $s_{\mathcal{I}} \in \Gamma(\mathcal{I}, \mathcal{L}|_{\mathcal{I}})$, which, by definition, is not the pull-back of a section of $\Gamma(\overline{\mathcal{M}}_{1,1}, \mathcal{L}) = 0$. That is to say, $s_{\mathcal{I}} \neq 0$. On the other hand if we take the trace $\text{tr}(s_{\mathcal{I}}) \in \Gamma(\overline{\mathcal{M}}_{1,1}, \mathcal{L}) = 0$ of $s_{\mathcal{I}}$, we get zero, so we conclude that $s_{\mathcal{I}}^2 \in \Gamma(\mathcal{I}, \mathcal{L}^2|_{\mathcal{I}})$ is the pull-back of a section of $\Gamma(\overline{\mathcal{M}}_{1,1}, \mathcal{L}^2)$, i.e., that $s_{\mathcal{I}}^2$ is a nonzero multiple of c_4 . Thus, $s_{\mathcal{I}}$ is of the form $\nu \cdot d_4$ (for some $\nu \in k$), and $\mathcal{I} \cong \mathcal{D}$. These observations, together with Lemma 1.1, prove the following:

Lemma 1.2. *$(\overline{\mathcal{N}}_{1,1}^\rho)_{\text{red}}$ has an irreducible component of degree exactly two over $\overline{\mathcal{M}}_{1,1}$ if and only if there exists a $\nu \in k^\times$ such that the two elements $\nu, -\nu \in k \cong \mathcal{L}|_\infty \subseteq \mathcal{B} = \overline{\mathcal{S}}_{1,1}^\rho$ (where the trivialization of $\mathcal{L}|_\infty$ is given by the residue map) lie in $\overline{\mathcal{N}}_{1,1}^\rho$.*

By Chapter I, Theorem 5.5, since $\deg(\lambda)$ (where λ is as in Chapter I, Theorem 5.5) is $\frac{1}{12}$ times the degree of the point at infinity, it follows that the canonical trivialization of $\overline{\mathcal{S}}_{1,1}^\rho|_\infty = \mathcal{L}_\infty \cong k$ is

$$\tau_\rho \stackrel{\text{def}}{=} \frac{1}{12} \left(\frac{1}{4} - \rho^2 \right) \equiv 2(1 + \rho^2) \pmod{5}$$

By Proposition 1.5 of Chapter II, the other two elements of $\overline{\mathcal{N}}_{1,1}^\rho|_\infty \subseteq \mathcal{L}_\infty \cong k$ are those that correspond to monodromy endomorphisms whose determinant is of the form $-r^2$ for some $r \in \mathbf{F}_5^\times$. Thus, by considering the *standard monodromy endomorphisms* of Chapter I, Definition 4.2, it follows that these other two elements of $\overline{\mathcal{N}}_{1,1}^\rho|_\infty \subseteq \mathcal{L}_\infty \cong k$ are those of the form $\tau_\rho + r^2$, for $r \in \mathbf{F}_5^\times$, i.e., $\tau_\rho \pm 1$. Now we compute the elements of $\overline{\mathcal{S}}_{1,1}^\rho|_\infty = \mathcal{L}_\infty \cong k$ that lie in $\overline{\mathcal{N}}_{1,1}^\rho$ case by case:

- (1) *the case $\rho = 0$:* We have $\{1, 2, 3\}$; the pair $\{2, 3\}$ satisfies the hypotheses of Lemma 1.2, and so we conclude that when $\rho = 0$, $(\overline{\mathcal{N}}_{1,1}^\rho)_{\text{red}}$ has exactly one component of degree two over $\overline{\mathcal{M}}_{1,1}$; this component passes through the elements 2 and 3 at infinity, and is étale over $\overline{\mathcal{M}}_{1,1}$ at the element 2. Note that since there do not exist any other degree two irreducible components, and,

moreover, there do not exist any degree one irreducible components (since $0 \notin \{1, 2, 3\}$), it follows that $\overline{\mathcal{N}}_{1,1}$ is reduced, and that the only other irreducible component of $\overline{\mathcal{N}}_{1,1}$ is of degree three over $\overline{\mathcal{M}}_{1,1}$. On the other hand, note that by Propositions 1.2 and 1.4 of Chapter II, every indigenous bundle in $\overline{\mathcal{N}}_{1,1}$ is admissible, so by [Mzk1], Chapter II, Corollary 2.16, $\mathcal{N}_{1,1}$ is smooth over the base field, hence reduced. Since $\overline{\mathcal{N}}_{1,1}$ is flat over $\overline{\mathcal{M}}_{1,1}$, it thus follows that $\overline{\mathcal{N}}_{1,1}$ is *reduced*. In other words, our calculation is compatible with what we know from the general theory.

- (2) *the case $\rho^2 = 1$:* we have $\{0, -1, -2\}$; thus, there are no pairs satisfying the hypotheses of Lemma 1.2. Since $0 \in \{0, -1, -2\} = \overline{\mathcal{N}}_{1,1}^1|_\infty$, it follows that there exists an irreducible component of $(\overline{\mathcal{N}}_{1,1}^1)_{\text{red}}$, which we shall denote \mathcal{I}_0 , which passes through 0 and is the unique irreducible component of $(\overline{\mathcal{N}}_{1,1}^1)_{\text{red}}$ of degree one over $\overline{\mathcal{M}}_{1,1}$. On the other hand, it follows *either* from the explicit construction of Lemma 1.3 below *or* (in fact) from the general theory of dormant bundles (cf. §2 of Chapter II, and §2 of this Chapter) that there exists an irreducible component \mathcal{I}_d of $(\overline{\mathcal{N}}_{1,1}^1)_{\text{red}}$ over which the tautological torally indigenous bundle is dormant. Let $\mathcal{I}_n \subseteq \overline{\mathcal{S}}_{1,1}^1$ be the first infinitesimal neighborhood of \mathcal{I}_d in $\overline{\mathcal{S}}_{1,1}^1$ (i.e., the subscheme obtained defined by the square of the ideal defining \mathcal{I}_d). Since the p -curvature of the tautological torally indigenous bundle on \mathcal{I}_d *vanishes identically* on \mathcal{I}_d , its *determinant automatically vanishes* on \mathcal{I}_n . Thus, $\mathcal{I}_n \subseteq \overline{\mathcal{N}}_{1,1}^1$, and so it follows that two times the degree of \mathcal{I}_d over $\overline{\mathcal{M}}_{1,1}$ is ≤ 5 . Thus, $\deg(\mathcal{I}_d/\overline{\mathcal{M}}_{1,1})$ is 1 or 2; but we know (by Lemma 1.2) that it cannot be 2. Hence, it follows that $\mathcal{I}_d = \mathcal{I}_0$. Also, it follows that $(\overline{\mathcal{N}}_{1,1}^1)_{\text{red}}$ has exactly one other irreducible component, of degree 3 over $\overline{\mathcal{M}}_{1,1}$, at which $\overline{\mathcal{N}}_{1,1}^1$ is necessarily reduced.

- (3) *the case $\rho^2 = -1$:* we have $\{1, 0, -1\}$; thus, the pair $\{1, -1\}$ satisfies the hypotheses of Lemma 1.2. Note also, that when $\rho = \pm \frac{1}{2}$, we have $\rho^2 = -1$, so that we are in the situation described immediately after Chapter I, Theorem 5.5. There, a canonical isomorphism was given between $\overline{\mathcal{S}}_{1,1}^{\frac{1}{2}}$ and $\overline{\mathcal{S}}_{1,0}$. It is easy to see that this isomorphism induces an isomorphism between $\overline{\mathcal{N}}_{1,1}^{\frac{1}{2}}$ and $\overline{\mathcal{N}}_{1,0}$. On the other hand, $\overline{\mathcal{N}}_{1,0}$ is described in detail in

[Mzk1], Chapter II, Corollary 3.10. In particular, this description is consistent with the conclusion obtained from Lemma 1.2, that $(\overline{\mathcal{N}}_{1,1}^{\frac{1}{2}})_{\text{red}}$ has an irreducible component of degree exactly two over $\overline{\mathcal{M}}_{1,1}$.

This completes the description of the various cases, modulo Lemma 1.3.

§1.2. The Case of Radius 1

Lemma 1.3. $\overline{\mathcal{N}}_{1,1}^1$ has a one-dimensional subscheme over which the tautological torally indigenous bundle is dormant.

Proof. Let S be a smooth, one-dimensional scheme over k . Let $(f : E \rightarrow S, \tau_4 : S \rightarrow E)$ be a one-pointed smooth curve of genus one over S . We shall write E^{\log} for the log scheme whose underlying scheme is E and whose log structure is that determined by the divisor $\text{Im}(\tau_4)$. Let us suppose that, when we take τ_4 as the origin of S , the remaining 2-torsion points $\tau_1, \tau_2, \tau_3 : S \rightarrow E$ are all defined over S . Let $P \stackrel{\text{def}}{=} \mathbf{P}(f_* \mathcal{O}_E(2\tau_4)) \cong \mathbf{P}_S^1$ be the trivialization such that the corresponding morphism $\phi : E \rightarrow P \cong \mathbf{P}_S^1$ (defined by the linear system P) maps τ_1 (respectively, τ_2, τ_3) to 0 (respectively, 1, ∞). We shall henceforth identify P with \mathbf{P}_S^1 , and denote the point 0 (respectively, 1, ∞) by σ_1 (respectively, σ_2, σ_3). Next, let us assume that we are given a finite, flat morphism $\psi : P \rightarrow Q \stackrel{\text{def}}{=} \mathbf{P}_S^1$ of degree four which is étale away from $\sigma_1, \sigma_2, \sigma_3$, and which has ramification index 3 at $\sigma_1, \sigma_2, \sigma_3$. It is easy to see that there exist data satisfying the above conditions: Indeed, since we are in characteristic 5, one may take for ψ the covering of $Q = \mathbf{P}_S^1$ which is ramified over three points of Q with ramification given by the permutations (123), (134), (142) on four letters. (Note that the product of these three permutations in the order given is equal to the identity.)

Let $(\mathcal{E}, \nabla_{\mathcal{E}})$ be the trivial vector bundle of rank two on $Q = \mathbf{P}_S^1$, endowed with the trivial connection. Let $\mathcal{L} \subseteq \mathcal{E}$ be the kernel of the tautological quotient (obtained by thinking of \mathbf{P}_S^1 as the moduli space of quotients of \mathcal{E}). Note that the Kodaira-Spencer map of $\mathcal{L} \subseteq \mathcal{E}$ is an isomorphism. Let $(\mathcal{F}, \nabla_{\mathcal{F}}) \stackrel{\text{def}}{=} (\psi \circ \phi)^*(\mathcal{E}, \nabla_{\mathcal{E}})$; $\mathcal{K} \stackrel{\text{def}}{=} (\psi \circ \phi)^*\mathcal{L}$. We would like to construct a new vector bundle with connection $(\mathcal{G}, \nabla_{\mathcal{G}})$ on E^{\log} out of $(\mathcal{F}, \nabla_{\mathcal{F}})$. This new bundle \mathcal{G} will be the same as \mathcal{F} except at τ_1, \dots, τ_4 , and the connection $\nabla_{\mathcal{G}}$ will have a logarithmic pole at τ_4 . At τ_i , for $i = 1, \dots, 4$, let $x \in \Gamma(U, \mathcal{O}_E)$ (where $U \subseteq E$ is open; $\text{Im}(\tau_i) \subseteq U$; $\text{Im}(\tau_j) \cap U = \emptyset$ for $j \neq i$) be a local section of \mathcal{O}_E such that $V(x) = \text{Im}(\tau_i)$. Then, we take $\mathcal{G}|_U$ to be the direct limit of the following diagram:

$$\begin{array}{ccc}
 \mathcal{K}|_U & \longrightarrow & \mathcal{F}|_U \\
 \downarrow & & \\
 \mathcal{K}|_U & &
 \end{array}$$

where the horizontal arrow is the natural inclusion (pulled back from Q), and the vertical arrow is multiplication by x^5 (respectively, x^2) if $i = 1, 2, 3$ (respectively, $i = 4$). One checks easily that when $i = 1, 2, 3$ (respectively, $i = 4$), $\nabla_{\mathcal{F}}$ stabilizes $\mathcal{G}|_U$ (respectively, has a logarithmic pole), hence induces a connection $\nabla_{\mathcal{G}}|_U$ on $\mathcal{G}|_U$ of the type desired. This completes the definition of $(\mathcal{G}, \nabla_{\mathcal{G}})$. Let $\mathcal{H} \subseteq \mathcal{G}$ be the subbundle induced by $\mathcal{K} \subseteq \mathcal{F}$. Then it is easy to check that the Kodaira-Spencer morphism of $\nabla_{\mathcal{G}}$ (as a logarithmic connection on E^{\log}) at $\mathcal{H} \subseteq \mathcal{G}$ is an isomorphism. (Indeed, this is clear away from τ_1, \dots, τ_4 . At τ_1, τ_2, τ_3 , the Kodaira-Spencer morphism for $(\mathcal{L} \hookrightarrow \mathcal{E})|_E$ has a zero of order 5 resulting from the fact that $E \rightarrow Q$ has ramification index $2 \cdot 3 = 6$; but this zero of order 5 is canceled by the above adjustment of the integral structure by multiplication by x^5 at τ_1, τ_2, τ_3 . Similarly, the above adjustment of the integral structure at τ_4 (combined with the fact that the ramification index of $E \rightarrow Q$ at τ_4 is 2) implies that the Kodaira-Spencer morphism is an isomorphism at τ_4 , as well.) It thus follows that $\mathbf{P}(\mathcal{G}, \nabla_{\mathcal{G}})$ is a torally indigenous bundle of radius ± 1 . Finally, since we constructed $(\mathcal{G}, \nabla_{\mathcal{G}})$ from the trivial bundle with connection $(\mathcal{E}, \nabla_{\mathcal{E}})$, it is clear that the p -curvature of $\mathbf{P}(\mathcal{G}, \nabla_{\mathcal{G}})$ vanishes identically. This completes the proof. \circ

§1.3. Conclusions

We summarize our conclusions in the following:

Theorem 1.4. *When $p = 5$, the structure of $\overline{\mathcal{N}}_{1,1}^{\rho}$ is as follows:*

- (1) the case $\rho = 0$: $\overline{\mathcal{N}}_{1,1}^0$ is reduced, and has one irreducible component of degree two over $\overline{\mathcal{M}}_{1,1}$, and one of degree three over $\overline{\mathcal{M}}_{1,1}$;
- (2) the case $\rho^2 = 1$: $(\overline{\mathcal{N}}_{1,1}^1)_{\text{red}}$ has one irreducible component of degree one over $\overline{\mathcal{M}}_{1,1}$ at which $\overline{\mathcal{N}}_{1,1}^1$ is nonreduced, and one irreducible component of degree three over $\overline{\mathcal{M}}_{1,1}$ at which $\overline{\mathcal{N}}_{1,1}^1$ is reduced;
- (3) the case $\rho^2 = -1$: $\overline{\mathcal{N}}_{1,1}^{-1} \cong \overline{\mathcal{N}}_{1,0}$ (see [Mzk1], Chapter II, Corollary 3.10 for more details).

Remark. Of course, this counterexample to the conjecture that $\overline{\mathcal{N}}_{g,r}$ is irreducible is vulnerable to the criticism that p is too specific: i.e., perhaps for large p , $\overline{\mathcal{N}}_{1,1}$ is irreducible. Nonetheless, even though, at the time of writing, I don't know how to construct counterexamples for arbitrarily large p , one aspect of the prime 5 – unlike 2 or 3 – is that it is not (at least in any obvious fashion) “characteristic to the situation.”

Remark. One interesting conclusion that one can draw from this calculation is the following: Consider the irreducible component I of $\overline{\mathcal{N}}_{1,1}$ which is of degree 3 over $\overline{\mathcal{M}}_{1,1}$. The discriminant of I over $\overline{\mathcal{M}}_{1,1}$ is easily seen to be a modular form of weight 12 (i.e., a section of $\mathcal{L}^{\otimes 6}$) on $\overline{\mathcal{M}}_{1,1}$ which is not identically zero, but has a zero at infinity. It thus follows that the discriminant of I over $\overline{\mathcal{M}}_{1,1}$ is nonzero over $\mathcal{M}_{1,1}$, i.e., that I is étale over $\mathcal{M}_{1,1}$. In particular, it follows that the ordinary locus of $\overline{\mathcal{N}}_{1,1}$ surjects onto $\overline{\mathcal{M}}_{1,1}$. That is to say, *at least for $p = 5$, every $g = 1, r = 1$ curve is (hyperbolically) ordinary* (in the sense of [Mzk1], Chapter II, Definition 3.3)! Needless to say, this behavior is quite different from the case $g = 1, r = 0$, where nonordinary curves exist for every prime. It is thus tempting to conjecture (see Introduction, §2.1, Question (1)) that, more generally, for all $p \geq 3$, and all g, r such that $2g - 2 + r \geq 1$, $\overline{\mathcal{N}}_{g,r}^{\text{ord}} \rightarrow \overline{\mathcal{M}}_{g,r}$ is surjective. At the time of writing, the author does not know how to prove such a conjecture.

§2. Higher Order Connections and Lubin-Tate Stacks

In this §, we classify all higher order connections on the projective line minus three points, as well as on elliptic curves. Using this classification, we show the nonemptiness of “Lubin-Tate stacks” which carry the same information as a certain portion of the stacks parametrizing higher order connections that we saw in Chapter II, except in a more useful and natural form.

§2.1. The Projective Line Minus Three Points

In this subsection, we give a complete classification of n -connections in the case $(g, r) = (0, 3)$.

Let S^{\log} be a $\mathbf{Z}/p^{n+1}\mathbf{Z}$ -flat noetherian log scheme equipped with the trivial log structure; $f^{\log} : X^{\log} \rightarrow S^{\log}$ be the 3-pointed smooth log-curve of genus 0; and $(\pi : P \rightarrow X, \nabla_P)$ a dormant n -connection

of radii ρ on a torally indigenous bundle on X_0^{\log} . Thus, we may regard (P, ∇_P) as a torally indigenous bundle on X^{\log} . (Indeed, since X^{\log} has no nontrivial deformations, the Hodge section of $P_0 \rightarrow X_0$ automatically lifts to $P \rightarrow X$.) Note further that the isomorphism class of the torally indigenous bundle (P, ∇_P) is *completely determined by its radii ρ* (cf. Chapter I, Theorem 4.4). Thus, there is no loss of generality in assuming S_0 to be the spectrum of an algebraically closed field k . We shall denote the three marked points of X^{\log} by the symbols $[0], [1], [\infty]$. Recall that by Chapter II, Lemma 2.3, (1), it follows that the radii $\rho_i \in (\mathbf{Z}/p^{n+1}\mathbf{Z})^\times$ (where $i = [0], [1], [\infty]$); for each radius ρ_i , let us fix an integer λ_i such that $0 \leq \lambda_i \leq p^{n+1} - 1$ and $2\rho_i \equiv \lambda_i \pmod{p^{n+1}}$. Also, let us assume that *the sum of the three λ_i is odd*.

Let $U \subseteq X$ be the complement of the three marked points. Let

$$\mathcal{L} \stackrel{\text{def}}{=} \mathcal{O}_X(\lambda_0[0] + \lambda_1[1] + \lambda_\infty[\infty])$$

where we regard the marked points as divisors in X . Let $\nabla_{\mathcal{L}}$ be the unique logarithmic connection on \mathcal{L} whose restriction to U is the trivial connection on $\mathcal{L}|_U = \mathcal{O}_U$.

Given a \mathbf{P}^1 -bundle $Q \rightarrow Y$ over an S -proper smooth curve Y , let us call Q *odd* (respectively, *even*) if there does not exist (respectively, exists) a line bundle \mathcal{L} on Q such that $\mathcal{L}^{\otimes 2} \cong \tau_{Q/Y}$. Note that to say that Q is odd (respectively, even) is equivalent to saying that it may be written as the projectivization of a rank two vector bundle of odd (respectively, even) degree. Since in the case under consideration, $2g - 2 + r = 1$, it follows that the \mathbf{P}^1 -bundle $P \rightarrow X$ is *odd*. Since the degree of \mathcal{L} is also *odd*, it follows that there exists a line bundle \mathcal{T} on P such that $\mathcal{T}^{\otimes 2} \cong \tau_{P/X} \otimes_{\mathcal{O}_X} \mathcal{L}$. Let $\mathcal{E} = \pi_* \mathcal{T}$. Thus, \mathcal{E} is a rank two vector bundle on X such that $\mathbf{P}(\mathcal{E}) = P$. Note, moreover, that the connections ∇_P and $\nabla_{\mathcal{L}}$ induce a connection on the family of polarized varieties given by $(\pi : P \rightarrow X, \mathcal{T})$, hence on the vector bundle \mathcal{E} . Let us denote this connection by $\nabla_{\mathcal{E}}$. Moreover, since we have a horizontal isomorphism $\mathbf{S}^2 \mathcal{E} \cong \text{Ad}(P) \otimes_{\mathcal{O}_X} \mathcal{L}$, it follows (by taking the determinant of both sides) that $\det(\mathcal{E}, \nabla_{\mathcal{E}})^{\otimes 3} \cong (\mathcal{L}, \nabla_{\mathcal{L}})^{\otimes 3}$. Now if $p \neq 3$, then this implies immediately that

$$\det(\mathcal{E}, \nabla_{\mathcal{E}}) \cong (\mathcal{L}, \nabla_{\mathcal{L}})$$

In fact, even if $p = 3$, we still obtain such an isomorphism since everything lifts to \mathbf{Z}_3 .

Let $\mathcal{E}^{[0]} = \mathcal{E}_0$. Thus, $\mathcal{E}^{[0]}$ is equipped with a connection $\nabla_{\mathcal{E}^{[0]}}$. The fact that $(P \rightarrow X, \nabla_P)$ is dormant implies that the p -curvature of $\nabla_{\mathcal{E}^{[0]}}$ is zero. Next, note that the Hodge section of the torally indigenous bundle $P_0 \rightarrow X_0$ defines a filtration $\mathcal{H} \subseteq \mathcal{E}^{[0]}$, where \mathcal{H} and $\mathcal{E}^{[0]}/\mathcal{H}$ are line bundles. Now we would like to inductively define a rank two vector

bundle $\mathcal{E}^{[i]}$ on $X_0^{F^i}$ (for $i \leq n+1$) whose restriction $\mathcal{E}^{[i]}|_{U_0^{F^i}}$ is equipped with a dormant connection $\nabla_{\mathcal{E}^{[i]}}$ (for $i \leq n$), as follows: Let $\mathcal{E}^{[i+1]} \subseteq \mathcal{E}^{[i]}$ be the subsheaf of sections annihilated by the connection $\nabla_{\mathcal{E}^{[i]}}$. Then we define the connection $\nabla_{\mathcal{E}^{[i+1]}}$ on $\mathcal{E}^{[i+1]}|_{U_0^{F^{i+1}}}$ via the same procedure as in the discussion preceding Definition 2.2 of Chapter II (where we defined the connections $\nabla_{\mathcal{A}^{[i]}}$). Let $\mathbf{P}^{[i]} \rightarrow X_0^{F^i}$ be the \mathbf{P}^1 -bundle obtained by taking the projectivization of $\mathcal{E}^{[i]}$. Note that we have an inclusion $\iota : (\Phi_{X/S}^{n+1})^* \mathcal{E}^{[n+1]} \subseteq \mathcal{E}^{[0]}$ which is an isomorphism over U_0 . Let $\mathcal{K} = \mathcal{H} \cap \text{Im}(\iota) \subseteq \mathcal{H}$. Thus, \mathcal{K} is a line bundle on X_0 which is a subsheaf of \mathcal{H} . Now we have the following result:

Lemma 2.1. *The torsion sheaf \mathcal{H}/\mathcal{K} has degree λ_x at the marked point x , and moreover, the natural injection $\mathcal{H}/\mathcal{K} \hookrightarrow \text{Coker}(\iota)$ is an isomorphism. In particular, the \mathbf{P}^1 -bundle $\mathbf{P}^{[n+1]} \rightarrow X_0^{F^{n+1}}$ is even.*

Proof. Let us first consider what happens at a marked point $x \in \{[0], [1], [\infty]\}$ of X^{\log} . Let V be an affine neighborhood of x such that $V - \{x\} \subseteq U$. Let t be a local parameter on V which is zero at x . By Chapter II, Lemma 2.3, it follows that (after possibly replacing V by a smaller neighborhood of x), one can write $(\mathcal{E}, \nabla_{\mathcal{E}})|_V$ in the following form:

$$(t^{-a} \cdot \mathcal{O}_V) \oplus (t^{-b} \cdot \mathcal{O}_V)$$

(where a, b are integers such that $0 \leq b \leq a \leq p^{n+1} - 1$), equipped with the connection induced by the trivial connection on $\mathcal{O}_V \oplus \mathcal{O}_V$. Moreover, $2\rho_x \equiv \pm(a-b) \pmod{p^{n+1}}$.

Since $\det(\mathcal{E}, \nabla_{\mathcal{E}}) = (\mathcal{L}, \nabla_{\mathcal{L}})$, it follows that $a+b \equiv \lambda_x \pmod{p^{n+1}}$. Since $a+b \leq 2p^{n+1}$, it thus follows that $a+b$ is either λ_x or $p^{n+1} + \lambda_x$. We claim that $a+b = \lambda_x$. Indeed, suppose that $a+b = p^{n+1} + \lambda_x$. Since $0 \leq a-b \leq p^{n+1}$ and $2\rho_x \equiv \pm(a-b) \pmod{p^{n+1}}$, it follows that $a-b$ is either λ_x or $p^{n+1} - \lambda_x$. Now since $a+b$ and $a-b$ have the same parity, we conclude that $a-b = p^{n+1} - \lambda_x$. But since we are operating under the assumption that $a+b = p^{n+1} + \lambda_x$, it thus follows that $2a = 2p^{n+1}$, so $a = p^{n+1}$, which is absurd. This proves the claim.

Thus, $a+b = \lambda_x$, and $a-b = \lambda_x$. In particular, $a = \lambda_x$, and $b = 0$. Thus, we have

$$\mathcal{E}_0|_{V_0} = (t^{-a} \cdot \mathcal{O}_{V_0}) \oplus \mathcal{O}_{V_0}$$

and (one computes easily that) $\text{Im}(\iota)|_{V_0} = \mathcal{O}_{V_0} \oplus \mathcal{O}_{V_0}$. Moreover, relative to this decomposition, $\mathcal{H} \subseteq \mathcal{E}_0$ is not contained in the right-hand factor modulo t . (Indeed, this follows from the fact that the Kodaira-Spencer

morphism for $\mathcal{H} \subseteq \mathcal{E}_0$ is an isomorphism.) It thus follows that \mathcal{H}/\mathcal{K} has degree $a = \lambda_x$ at x . Since $\text{Coker}(\iota)$ has degree $a + b = \lambda_x$ at x , it thus follows that the inclusion $\mathcal{H}/\mathcal{K} \hookrightarrow \text{Coker}(\iota)$ is an isomorphism. Since the sum of the λ_x 's is odd, it thus follows that the degree of $\text{Coker}(\iota)$ is odd. This implies that the parity of $\mathbf{P}^{[n+1]} = \mathbf{P}(\mathcal{E}^{[n+1]})$ is different from that of $P \rightarrow X$. Since $P \rightarrow X$ is odd, it thus follows that $\mathbf{P}^{[n+1]}$ is even. \bigcirc

Lemma 2.2. *The \mathbf{P}^1 -bundle $\mathbf{P}^{[n+1]} \rightarrow X_0^{F^{n+1}}$ is trivial, i.e., isomorphic to $\mathbf{P}_{X_0^{F^{n+1}}}^1 \rightarrow X_0^{F^{n+1}}$.*

Proof. First, let us make some general observations concerning even \mathbf{P}^1 -bundles on \mathbf{P}_k^1 : Any even \mathbf{P}^1 -bundle on \mathbf{P}_k^1 may be written as $\mathbf{P}(\mathcal{O}_{\mathbf{P}_k^1}(a) \oplus \mathcal{O}_{\mathbf{P}_k^1}(-a))$ for a unique nonnegative integer a . We shall refer to a as the *order* of the \mathbf{P}^1 -bundle. Thus, the trivial \mathbf{P}^1 -bundle on \mathbf{P}_k^1 may be characterized as the unique even \mathbf{P}^1 -bundle of order 0. Next, let us make the following observation:

Any even \mathbf{P}^1 -bundle on \mathbf{P}_k^1 of nonzero order may be deformed to a \mathbf{P}^1 -bundle of strictly lower order.

Indeed, suppose that the \mathbf{P}^1 -bundle that we are given is the projectivization of $\mathcal{O}_{\mathbf{P}_k^1}(a) \oplus \mathcal{O}_{\mathbf{P}_k^1}(-a)$, for some positive integer a . Then the above observation follows from the fact that the projectivization of the *any* nontrivial extension of $\mathcal{O}_{\mathbf{P}_k^1}(a)$ by $\mathcal{O}_{\mathbf{P}_k^1}(-a)$ necessarily has order $< a$. (Note also that $H^1(\mathbf{P}_k^1, \mathcal{O}_{\mathbf{P}_k^1}(-2a))$ is nonzero if $a > 0$, so such nontrivial extensions always exist.) This completes the proof of the observation.

Now we would like to prove that the particular even \mathbf{P}^1 -bundle $\mathbf{P}^{[n+1]} \rightarrow X_0^{F^{n+1}}$ is trivial. By the discussion in the preceding paragraph, *it suffices to prove that this \mathbf{P}^1 -bundle is isomorphic to any deformation of itself.* We do this as follows. Let $T = \text{Spec}(A)$, where $A = (W(k)/p^{n+1}W(k))[[t]]$; k is an algebraically closed field; and t is an indeterminate. Thus, we have a natural morphism $T \rightarrow S = \text{Spec}(W(k)/p^{n+1}W(k))$. Let $Y^{\log} \rightarrow T^{\log}$ (where T^{\log} is T equipped with the trivial log structure) be the pull-back of $X^{\log} \rightarrow S^{\log}$ to T^{\log} . We would like to regard $X^{\log} \subseteq Y^{\log}$ and $S^{\log} \subseteq T^{\log}$ by means of the embeddings defined by $t = 0$. Let $\mathbf{P}_Y^{[n+1]} \rightarrow Y_0^{F^{n+1}}$ be *any* \mathbf{P}^1 -bundle whose restriction to $X_0^{F^{n+1}} \subseteq Y_0^{F^{n+1}}$ is $\mathbf{P}^{[n+1]} \rightarrow X_0^{F^{n+1}}$. Then we propose to show that $\mathbf{P}_Y^{[n+1]} \rightarrow Y_0^{F^{n+1}}$ is isomorphic to the pull-back via the natural projection $Y_0^{F^{n+1}} \rightarrow X_0^{F^{n+1}}$ of a \mathbf{P}^1 -bundle on $X_0^{F^{n+1}}$.

The crucial point is to observe that we may reverse the steps in the construction of $\mathbf{P}^{[n+1]}$, that is:

- (1) We lift $\mathbf{P}_Y^{[n+1]} \rightarrow Y_0^{F^{n+1}}$ to some \mathbf{P}^1 -bundle $\tilde{\mathbf{P}}_Y^{[n+1]} \rightarrow Y$.

- (2) If, Zariski locally on Y , we choose a connection for this \mathbf{P}^1 -bundle $\tilde{\mathbf{P}}_Y^{[n+1]} \rightarrow Y$, then we get a crystal in \mathbf{P}^1 -bundles which we may pull back by the relative Frobenius of Y_0 over T a total of $n+1$ times. Moreover, it is easy to see that by pulling back by Frobenius $n+1$ times, the resulting pulled back crystal is *independent* of the choice of connection. Thus, we get a natural crystal in \mathbf{P}^1 -bundles $(\bar{\mathbf{P}}_Y^{[n+1]} \rightarrow Y, \nabla_{\bar{\mathbf{P}}_Y^{[n+1]}})$ which is the result of pulling back $\tilde{\mathbf{P}}_Y^{[n+1]} \rightarrow Y$ by Frobenius $n+1$ times.
- (3) We adjust the integral structure of $\bar{\mathbf{P}}_Y^{[n+1]} \rightarrow Y$ at the marked points of Y^{\log} (i.e., if y is a marked point of Y^{\log} with local parameter t , and we write $\bar{\mathbf{P}}_Y^{[n+1]}$ locally as the projectivization of some rank two vector bundle, then we allow denominators dividing t^{λ_y} for sections of some horizontal rank one line bundle of this vector bundle) so that the resulting crystal in \mathbf{P}^1 -bundles (Q, ∇_Q) is of radii ρ . Note that this can be done in such a way that the restriction of (Q, ∇_Q) to $X_0^{\log} \subseteq Y_0^{\log}$ is (P_0, ∇_{P_0}) . (Indeed, this follows from the discussion of the explicit local structure of $(P, \nabla_P) = \mathbf{P}(\mathcal{E}, \nabla_{\mathcal{E}})$ in the proof of Lemma 2.1.)

Thus, when we apply to (Q, ∇_Q) the procedure carried out above to construct $\mathbf{P}^{[n+1]}$ out of (P, ∇_P) , we recover $\mathbf{P}_Y^{[n+1]}$. Note, moreover, that *since there are no obstructions to deforming the Hodge section* (i.e., $H^1(X_0, \tau_{X_0^{\log}/S_0^{\log}}) = 0$), it follows that the Hodge section of $Q|_{X_0} = P_0$ deforms to a section of $Q \rightarrow Y$, hence that (Q, ∇_Q) is totally indigenous on Y^{\log} of radii ρ . But since Y^{\log} admits only one such bundle (cf. Chapter I, Theorem 4.4) it follows that (Q, ∇_Q) is just the pull-back of (P, ∇_P) to Y^{\log} . In particular, $\mathbf{P}_Y^{[n+1]} \rightarrow Y_0^{F^{n+1}}$ is the pull-back to $Y_0^{F^{n+1}}$ of $\mathbf{P}^{[n+1]} \rightarrow X_0^{F^{n+1}}$, as desired. This completes the proof of the lemma. \circ

Let us review the situation obtained so far. First of all, let $\mathcal{F} \stackrel{\text{def}}{=} (\Phi_{X/S}^{n+1})^* \mathcal{E}^{[n+1]} \subseteq \mathcal{E}^{[0]}$. Let $\Lambda = \lambda_{[0]} + \lambda_{[1]} + \lambda_{[\infty]}$. Then $\deg(\mathcal{F}) = \deg(\mathcal{E}) - \Lambda = 0$ (by Lemma 2.1). Moreover, by Lemma 2.2, $Q_0 \stackrel{\text{def}}{=} \mathbf{P}(\mathcal{F}) \rightarrow X_0$ is the trivial \mathbf{P}^1 -bundle. Fix a trivialization $Q_0 \cong \mathbf{P}_{X_0}^1$. Also, we have a locally split injection $\mathcal{K} \hookrightarrow \mathcal{F}$ which determines a section $s_{\mathcal{K}} : X_0 \rightarrow Q_0$. Let $\psi : X_0 \rightarrow Q_0 \rightarrow \mathbf{P}_k^1$ be the composite of $s_{\mathcal{K}}$ with the projection $Q_0 \rightarrow \mathbf{P}_k^1$ arising from the trivialization.

Note that the Kodaira-Spencer morphism for $\mathcal{H} \subseteq \mathcal{E}^{[0]}$ defines an isomorphism of $\mathcal{H}^{\otimes 2} \otimes \mathcal{L}_0^{-1} \cong \omega_{X_0^{\log}/S_0^{\log}}$. This isomorphism tells us that

the derivative $d\psi : \psi^* \omega_{\mathbf{P}_k^1/k} \rightarrow \omega_{X_0/k}$ is an isomorphism away from the marked points of X_0^{\log} , i.e., that ψ is étale away from the marked points of X_0^{\log} . By taking degrees of both sides, this isomorphism also tells us that $2 \deg(\mathcal{H}) = \Lambda + 1$. On the other hand, by Lemma 2.1, $\deg(\mathcal{K}) + \Lambda = \deg(\mathcal{H})$. Thus, $2 \deg(\mathcal{K}) = 1 - \Lambda$. Since $\deg(\psi^* \omega_{\mathbf{P}_k^1/k}) = 2 \deg(\mathcal{K})$, we thus obtain that the zero locus of $d\psi$ has degree $-2 - 2 \deg(\mathcal{K}) = \Lambda - 3$.

Next, let us observe that by the explicit analysis of the proof Lemma 2.1, ψ has ramification index λ_x at the marked point x . Indeed, in the notation of the proof of Lemma 2.1, \mathcal{H} is not contained in either the left-hand or right-hand factors of $\mathcal{E}_0|_{V_0} = (t^{-\lambda_x} \cdot \mathcal{O}_{V_0}) \oplus \mathcal{O}_{V_0}$ modulo t . Thus, \mathcal{K} coincides with the left-hand factor of $\mathrm{Im}(\iota)|_{V_0} = \mathcal{F}|_{V_0} = \mathcal{O}_{V_0} \oplus \mathcal{O}_{V_0}$ to order precisely λ_x . But this is just another way of saying that the pull-back via ψ of a coordinate on \mathbf{P}_k^1 at $\psi(x)$ has a zero of order precisely λ_x at x . This completes the proof of the observation.

Thus, the morphism $d\psi : \psi^* \omega_{\mathbf{P}_k^1/k} \rightarrow \omega_{X_0/k}$ induced on differentials by ψ is an isomorphism over U_0 and has a zero of order exactly $\lambda_x - 1$ at the marked point x . This is consistent with the assertion (two paragraphs above) that the zero locus of $d\psi$ has degree $\Lambda - 3$. In summary,

$\psi : X_0 \rightarrow \mathbf{P}_k^1$ is a finite, separable, tamely ramified morphism between two copies of \mathbf{P}_k^1 that is tamely ramified at ≤ 3 points upstairs.

It follows immediately from the Riemann-Hurwitz formula that the images of the three marked points of X_0 in \mathbf{P}_k^1 are distinct.

On the other hand, suppose that we are given a finite, separable morphism $\psi_0 : X_0 \rightarrow \mathbf{P}_k^1$ that is tamely ramified at each marked point x of X_0 with ramification index λ_x , and étale everywhere else. For simplicity, let us assume that ψ_0 maps $[0]$ (respectively, $[1]$; $[\infty]$) to $[0]$ (respectively, $[1]$; $[\infty]$). Thus, we may regard ψ_0 as a morphism $X_0 \rightarrow X_0$ that takes marked points to marked points. Let $\psi : X \rightarrow X$ be the unique lifting of ψ_0 that maps marked points to marked points. (Note that such a ψ exists since ψ_0 is tamely ramified.) Let \mathcal{F} be the trivial bundle on X , and let $F^1(\mathcal{F}) \subseteq \mathcal{F}$ be the kernel of the tautological quotient (obtained by thinking of $X = \mathbf{P}_S^1$ as a moduli space of quotients of \mathcal{F}). Let \mathcal{G} be the rank two vector bundle on X that is equal to \mathcal{F} away from the marked points, and in a neighborhood of $x \in \{[0], [1], [\infty]\}$, is the push-out of the following diagram:

$$\begin{array}{ccc} F^1(\mathcal{F}) & \longrightarrow & \mathcal{F} \\ \downarrow t_x & & \\ F^1(\mathcal{F}) & & \end{array}$$

where the horizontal arrow is the natural inclusion, and the vertical arrow is multiplication by t_x , a local uniformizer at x that vanishes at x . Let $\nabla_{\mathcal{F}}$ be the trivial connection on \mathcal{F} . Then $\nabla_{\mathcal{F}}$ induces a connection $\nabla_{\mathcal{G}}$ on \mathcal{G} with logarithmic singularities at $\{[0], [1], [\infty]\}$. Let $(\mathcal{D}, \nabla_{\mathcal{D}})$ be the vector bundle with connection on X^{\log} that is equal to $\psi^*(\mathcal{F}, \nabla_{\mathcal{F}})$ away from the marked points of X_0^{\log} , and equal to $\psi^*(\mathcal{G}, \nabla_{\mathcal{G}})$ in a neighborhood of the marked points of X^{\log} . Then $(Q, \nabla_Q) \stackrel{\text{def}}{=} \mathbf{P}(\mathcal{D}, \nabla_{\mathcal{D}}) \rightarrow X$ is a \mathbf{P}^1 -bundle over X equipped with a section $s_Q : X \rightarrow Q$ arising from $F^1(\mathcal{F}) \subseteq \mathcal{F}$. Moreover, one sees easily that the Kodaira-Spencer morphism for this section s_Q is an isomorphism. Thus, (Q, ∇_Q) is a dormant n -connection of radii ρ on the torally indigenous bundle (Q_0, ∇_{Q_0}) on X_0^{\log} . In summary, we see that we have proven the following

Theorem 2.3. *Let k be an algebraically closed field of odd characteristic p . Let $\rho_{[0]}, \rho_{[1]}, \rho_{[\infty]} \in (\mathbf{Z}/p^{n+1}\mathbf{Z})^\times$. For each radius $\rho_i \in (\mathbf{Z}/p^{n+1}\mathbf{Z})^\times$ (where $i = [0], [1], [\infty]$), let λ_i be an integer such that $0 \leq \lambda_i \leq p^{n+1} - 1$ and $2\rho_i \equiv \lambda_i \pmod{p^{n+1}}$. Let us assume, moreover, that the sum of the three λ_i 's is odd. Then there is a natural bijective correspondence between the following:*

(*) *dormant n -connections of radii ρ on the torally indigenous bundle of radii $\rho \pmod{p}$ on the 3-pointed curve of genus 0 over k*

and

(†) *finite, separable morphisms $\psi_0 : \mathbf{P}_k^1 \rightarrow \mathbf{P}_k^1$ that are ramified at $[0]$ (respectively, $[1]; [\infty]$) with index $\lambda_{[0]}$ (respectively, $\lambda_{[1]}; \lambda_{[\infty]}$) and étale elsewhere.*

Remark. Thus, we obtain a complete classification of dormant n -connections on the 3-pointed curve of genus 0. Of course, going from ψ_0 to a dormant n -connection is essentially trivial. Thus, the non-trivial part of the Theorem lies in the statement that all dormant n -connections are obtained in this way. The utility of this nontrivial part of the Theorem may be seen in the following observation: By using this part of the Theorem, one can show the existence of n -connections whose p^{n+1} -curvature is nilpotent, but *not* identically zero. For a concrete example of this, consider the following: Suppose that $p \geq 5$, and $n \geq 1$; then we take

$$\lambda_{[0]} = 2; \quad \lambda_{[1]} = 2; \quad \lambda_{[\infty]} = 3 + 2p^n$$

Then it follows from Theorem 2.3 that this triple of λ 's defines a *nondormant n -connection*. Indeed, the fact that this n -connection is nondormant follows by observing that Theorem 2.3 implies that if it were

dormant, then lifting the corresponding ψ_0 to characteristic zero and using the complex analytic description of the fundamental group of \mathbf{P}^1 minus three points, we would obtain 2 transpositions whose product is a cyclic permutation of order $3+2p^n$, which is absurd. Moreover, the nilpotence of the n -connection is automatic, since the n -Verschiebung is a section of a vector bundle of rank $3g-3+r=0$. The reader may amuse himself (respectively, herself) by constructing similar concrete examples of n -connections by means of Theorem 2.3 and the complex analytic description of the fundamental group of \mathbf{P}^1 minus three points.

Remark. Suppose that $n=0$. Then for a given λ_0, λ_1 , and λ_∞ , it is always the case that either $\lambda_0+\lambda_1+\lambda_\infty < 2p$, or $(p-\lambda_0)+(p-\lambda_1)+\lambda_\infty < 2p$. (Indeed, if this were not so, then it would follow that $\lambda_\infty \geq p$, which is absurd.) Thus, we may always choose the λ_i so that their sum Λ is $< 2p$. The significance of this observation is as follows: By the Riemann-Hurwitz formula, $2 \deg(\psi) = \Lambda - 1$ (where $\psi = \psi_0$ is as in (\dagger)), so we obtain that $\deg(\psi) \leq p-1$. But this means that a ψ as in (\dagger) exists in characteristic p if and only if it exists in characteristic 0. As noted in the preceding Remark, the complex analytic description of the fundamental group of \mathbf{P}^1 minus three points means that whether or not such a ψ exists in characteristic 0 is an *entirely combinatorial issue*. Thus, in summary, we see that

Theorem 2.3 gives an entirely combinatorial criterion for determining whether a characteristic p torally indigenous bundle on X^{\log} of given radii ρ is dormant or not.

Remark. By coupling the concrete examples one can construct from Theorem 2.3 with Theorem 2.8 of Chapter II, one obtains easily the result that there *exist* dormant indigenous bundles on smooth curves of genus g , for every $g \geq 2$. In particular, one thus obtains the existence of stable vector bundles (given, for instance, by $\mathrm{Ad}(P)^{\mathrm{hor}}$) on smooth curves of genus g (for every $g \geq 2$) whose Frobenius pull-back is indigenous, hence not semi-stable. I believe that the first examples of such stable bundles were constructed by Gieseker ([Gie]) over Mumford curves. Thus, we obtain an alternate proof of the existence of such stable bundles.

§2.2. Elliptic Curves

In this subsection, we perform an analysis similar to that of the preceding section, but this time for elliptic curves. Let $n \geq 0$. Let k be an algebraically closed field of characteristic p ; $S \stackrel{\mathrm{def}}{=} \mathrm{Spec}(W(k)/p^{n+1} \cdot W(k))$. Let S^{\log} be the log scheme obtained by equipping S with the trivial log structure; $f^{\log} : X^{\log} \rightarrow S^{\log}$ be a 1-pointed smooth log-curve

of genus 1; and $(\pi : P \rightarrow X, \nabla_P)$ an n -connection of radius $\rho = \frac{1}{2}$ on a torally indigenous bundle on X_0^{\log} . Let us assume that $X^{\log} \rightarrow S^{\log}$ has been chosen so that (P, ∇_P) defines a torally indigenous bundle on X^{\log} . By the technique discussed after Theorem 5.5 of Chapter I,

To study (P, ∇_P) is equivalent to studying (torally) indigenous bundles on the elliptic curve (i.e., 0-pointed curve of genus 1) X .

(Of course, when we say “(torally) indigenous bundles on elliptic curves,” the reader may object that so far we have only defined and studied such bundles under the assumption that $2g - 2 + r \geq 1$. However, it is essentially trivial to extend this theory to the case $g = 1$, $r = 0$.) Thus, let (Q, ∇_Q) be the indigenous bundle on X corresponding to (P, ∇_P) . Since X has no marked points or nodes, we can apply the discussion preceding Definition 2.2 of Chapter II (involving $\mathcal{A}^{[i]}$ and $\nabla_{\mathcal{A}^{[i]}}$) directly to X (that is, without restricting to some open subset $U \subseteq X$). Since, for X , $2g - 2 + r = 0$ is even, it follows that we can write (Q, ∇_Q) as the projectivization of a rank two bundle with connection $(\mathcal{E}, \nabla_{\mathcal{E}})$ whose determinant is trivial. Now, we let $(\mathcal{E}^{[0]}, \nabla_{\mathcal{E}^{[0]}}) = (\mathcal{E}, \nabla_{\mathcal{E}})_0$, and define $(\mathcal{E}^{[i]}, \nabla_{\mathcal{E}^{[i]}})$ ($0 \leq i \leq n$) inductively as in the discussion preceding Definition 2.2 of Chapter II. Thus, $\mathcal{E}^{[0]} = (\Phi_{X/S}^i)^* \mathcal{E}^{[i]}$. In particular, it follows that, for all $0 \leq i \leq n$, $\mathcal{E}^{[i]}$ is a *semi-stable vector bundle* on $X_0^{F^i}$. Note that the fact that the Kodaira-Spencer morphism associated to the Hodge filtration on \mathcal{E} is an isomorphism implies that \mathcal{E} is a decomposable bundle. (Indeed, the direct sum of the Hodge subbundle and “its derivative” gives a direct sum decomposition of \mathcal{E} .) Since we are ultimately only interested in the associated projective bundles, by tensoring \mathcal{E} with a line bundle whose square is trivial, we may thus assume henceforth that \mathcal{E} is, in fact, the trivial vector bundle of rank two.

Suppose that X_0 is an *ordinary elliptic curve* (i.e., its classical Hasse invariant is nonzero). We claim that for $i \geq 1$, $\mathcal{E}^{[i]}$ is then necessarily *decomposable*. Indeed, if this were not the case, then by [Harts1], Chapter V, §2, Theorem 2.15, it would follow that there exists a degree zero line bundle \mathcal{L} on $X_0^{F^i}$ such that we have an exact sequence

$$0 \rightarrow \mathcal{L}^{-1} \rightarrow \mathcal{E}^{[i]} \rightarrow \mathcal{L} \rightarrow 0$$

If this sequence is to be nonsplit, then $\mathcal{L}^{\otimes 2}$ must be trivial. Thus, it follows from the fact that X_0 is ordinary that the pull-back of this sequence via $\Phi_{X/S}^i$ is again nonsplit. But this contradicts the fact that $\mathcal{E}^{[0]}$ is trivial. Thus, $\mathcal{E}^{[i]}$ is decomposable.

Let us write $\mathcal{E}^{[i]} = \mathcal{L}^{[i]} \oplus (\mathcal{L}^{[i]})^{-1}$, where $\mathcal{L}^{[i]}$ is a line bundle of degree zero on $X_0^{F^i}$. Since $(\Phi_{X/S}^i)^* \mathcal{E}^{[i]}$ is trivial, it thus follows that $(\Phi_{X/S}^i)^* \mathcal{L}^{[i]}$ is trivial, hence that $(\mathcal{L}^{[i]})^{\otimes p^i} \cong \mathcal{O}_{X_0^{F^i}}$. We claim that $\mathcal{L}^{[1]}$ is nontrivial.

Indeed, if $\mathcal{L}^{[1]}$ were trivial, then $\nabla_{\mathcal{E}^{[0]}}$ would be the trivial connection on the trivial bundle $\mathcal{E}^{[0]}$, so $(\mathcal{E}^{[0]}, \nabla_{\mathcal{E}^{[0]}})$ could not be indigenous (that is to say, the trivial bundle (equipped with the trivial connection) does not contain a degree zero subbundle whose Kodaira-Spencer morphism is an isomorphism). This proves the claim. Thus, $\mathcal{L}^{[i]}$ is a generator of the set of line bundles on $X_0^{F^i}$ whose $(p^i)^{\text{th}}$ power is trivial.

Finally, note that if $n \geq 1$, then the p^{n+1} -curvature of (P, ∇_P) , which is just the p -curvature of $(\mathcal{E}^{[n]}, \nabla_{\mathcal{E}^{[n]}})$, defines (after choosing a trivialization of $\omega_{X_0^{F^n}/S_0}$) a section of $(\mathcal{L}^{[n]})^2 \oplus \mathcal{O}_{X_0^{F^n}} \oplus (\mathcal{L}^{[n]})^{-2}$ over $X_0^{F^n}$. Since $(\mathcal{L}^{[n]})^2$ is nontrivial (under the assumption $n \geq 1$), it thus follows that the p^{n+1} -curvature must lie in the middle factor $\mathcal{O}_{X_0^{F^n}}$. Thus, if the p^{n+1} -curvature is assumed to be nilpotent, we conclude that it must, in fact, be zero. That is to say, we have proven the following result:

Proposition 2.4. *Let $n \geq 0$. Let $S \stackrel{\text{def}}{=} \text{Spec}(W(k)/p^{n+1} \cdot W(k))$, where k is an algebraically closed field of characteristic p . Let $X \rightarrow S$ be an elliptic curve such that $X_0 \rightarrow S_0$ is ordinary. Then:*

- (1) *Suppose that $(P \rightarrow X, \nabla_P)$ is a dormant n -connection on an indigenous bundle on X_0 . Let $\mathcal{A}^{[i]}$ (for $i = 0, \dots, n+1$) be as in the discussion preceding Definition 2.2 of Chapter II. Then $\mathcal{A}^{[i]}$ is the “Ad” associated to a vector bundle of rank two of the form $\mathcal{L} \oplus \mathcal{L}^{-1}$ for some $\mathcal{L} \in \text{Pic}(X_0^{F^i}/S_0)$ which is annihilated by p^i , but not by p^{i-1} .*
- (2) *Let $(\mathcal{F}, \nabla_{\mathcal{F}})$ be a crystal in rank two vector bundles with trivial determinant on $\text{Crys}(X_0^{F^{n+1}}/S)$ such that $\mathcal{F}_0 = \mathcal{L} \oplus \mathcal{L}^{-1}$ for some $\mathcal{L} \in \text{Pic}(X_0^{F^{n+1}}/S_0)$ which is annihilated by p^{n+1} , but not by p^n . Then the pull-back $(\Phi_{X/S}^{n+1})^*(\mathcal{F}, \nabla_{\mathcal{F}})$ (that takes crystals on $\text{Crys}(X_0^{F^{n+1}}/S)$ to crystals on $\text{Crys}(X/S)$) is a dormant n -connection on an indigenous bundle on X_0 .*
- (3) *If $n \geq 1$, then every nilpotent n -connection on an indigenous bundle on X_0 is automatically dormant.*

Proof. (1) and (3) follow from the above discussion. To complete the proof of (2), it suffices to check that the reduction modulo p of such a $(\Phi_{X/S}^{n+1})^*(\mathcal{F}, \nabla_{\mathcal{F}})$ is always indigenous. But this follows from taking for the Hodge filtration the graph in $(\Phi_{X/S}^{n+1})^*\mathcal{L} \oplus (\Phi_{X/S}^{n+1})^*\mathcal{L}^{-1}$ of an isomorphism $(\Phi_{X/S}^{n+1})^*\mathcal{L} \cong \mathcal{O}_{X_0} \cong (\Phi_{X/S}^{n+1})^*\mathcal{L}^{-1}$, which, of course, can never be horizontal. This completes the proof of (2). \circ

Next, we consider the supersingular case. Thus, let X_0 be supersingular. Let us consider $\mathcal{E}^{[i]}$, for $i \geq 1$. If such an $\mathcal{E}^{[i]}$ were *decomposable*, then, by the same argument as in the ordinary case, it would have to be of the form $\mathcal{L} \oplus \mathcal{L}^{-1}$ for some *nontrivial* line bundle \mathcal{L} on $X_0^{F^i}$ such that $\mathcal{L}^{\otimes p^i} \cong \mathcal{O}_{X_0^{F^i}}$. But since X_0 is supersingular, such a line bundle \mathcal{L} cannot exist. Thus, it follows that for $i \geq 1$, $\mathcal{E}^{[i]}$ is *indecomposable*. In particular, each $\mathcal{E}^{[i]}$ must fit into a nonsplit exact sequence of the form $0 \rightarrow \mathcal{O}_{X_0^{F^i}} \rightarrow \mathcal{E}^{[i]} \rightarrow \mathcal{O}_{X_0^{F^i}} \rightarrow 0$. (Here, we use the facts that \mathcal{E} is trivial and that (since X_0 is supersingular) any line bundle on $X_0^{F^i}$ whose $(p^i)^{\text{th}}$ power is trivial is itself trivial.) But since this exact sequence splits after pull-back by Frobenius (since X_0 is supersingular), and $\mathcal{E}^{[i]} = \Phi_{X/S}^* \mathcal{E}^{[i+1]}$, we thus obtain that $\mathcal{E}^{[i]}$ cannot exist if $i \geq 2$. Thus, we obtain that $n \leq 1$, and that no 1-connection can be dormant. In other words, we have the following result:

Proposition 2.5. *Let $S \stackrel{\text{def}}{=} \text{Spec}(W(k)/p^{n+1} \cdot W(k))$, where k is an algebraically closed field of characteristic p . Let $X \rightarrow S$ be an elliptic curve such that $X_0 \rightarrow S_0$ is supersingular. Then:*

- (1) *Every nilpotent indigenous bundle on X_0 is dormant, and isomorphic to the pullback by $\Phi_{X/S}$ of the unique rank two vector bundle \mathcal{G} on X_0^F which is a nontrivial extension of $\mathcal{O}_{X_0^F}$ by $\mathcal{O}_{X_0^F}$. Moreover, $\Phi_{X/S}^* \mathcal{G}$ is a dormant indigenous bundle on X_0 .*
- (2) *Every 1-connection on an indigenous bundle on X_0 is necessarily nilpotent and nondormant.*
- (3) *There do not exist any n -connections on indigenous bundles on X_0 if $n \geq 2$.*

Proof. The first part of (1) follows from the explicit computations of the p -curvature in the discussion preceding [Mzk1], Chapter II, Theorem 3.9. The only thing left to prove is that $\Phi_{X/S}^* \mathcal{G}$ is indigenous. But this follows from taking for the Hodge filtration a (necessarily nonhorizontal) splitting of the extension sequence $0 \rightarrow \mathcal{O}_{X_0^F} \rightarrow \Phi_{X/S}^* \mathcal{G} \rightarrow \mathcal{O}_{X_0^F} \rightarrow 0$ obtained by pulling back the given extension sequence of \mathcal{G} . \square

§2.3. Lubin-Tate Stacks

Let $\varpi \geq 1$ be a natural number. Let Π_ϖ be the unique VF-pattern of period ϖ that assigns to every integer not divisible by ϖ the level 0.

Definition 2.6. We shall call Π_ϖ the *VF-pattern of pure tone ϖ* . We shall call the VF-stack (respectively, shifted VF-stack) associated to Π_ϖ the *VF-stack (respectively, shifted VF-stack) of pure tone ϖ* . (Recall from Definition 1.1 of Chapter III that) when $\varpi = 1$, we shall call Π_ϖ the *home VF-pattern*, and its associated VF-stack the *home VF-stack*.

Proposition 2.7. *Every VF-stack, as well as every shifted VF-stack, of pure tone $\varpi \geq 1$ is smooth of dimension $3g - 3 + r$ over \mathbf{F}_p .*

Proof. For VF-stacks, this follows from Proposition 1.7 and Lemma 1.8 of Chapter III. For shifted VF-stacks, this follows by the same argument (involving Proposition 1.5 of Chapter III) as that used in the proof of Corollary 1.6 of Chapter III. \circ

Let us set $\Pi = \Pi_\varpi$, and consider the natural morphism

$$\nu : \overline{\mathcal{N}}_{g,r}^{\Pi,s} \rightarrow \overline{\mathcal{M}}_{g,r}$$

It follows from the definition of $\overline{\mathcal{N}}_{g,r}^{\Pi,s}$ that if $\varpi \geq 2$, then $r = 0$, and the image of ν is contained in \mathcal{M}_g , the locus of smooth curves of genus g . We shall mainly be concerned in this subsection with the case $\varpi \geq 2$. Let us denote by $\mathcal{L}_{g,r}^\varpi \subseteq \overline{\mathcal{N}}_{g,r}^{\Pi,s}$ the open substack over which ν is étale.

Definition 2.8. We shall call $\mathcal{L}_{g,r}^\varpi$ the *(r -pointed) Lubin-Tate stack of (genus g and) tone ϖ* .

Thus, when $\varpi = 1$, we recover the stack $\overline{\mathcal{N}}_{g,r}^{\text{ord}}$ of [Mzk1]. When $\varpi \geq 2$, the stack \mathcal{L}_g^ϖ parametrizes “Lubin-Tate uniformizations” of a curve. We will discuss such uniformizations in more detail in Chapter VIII, but roughly speaking, just as the ordinary uniformizations of [Mzk1] induce local isomorphisms of (the ordinary locus of) the curve with \mathbf{G}_m , Lubin-Tate uniformizations induce local isomorphisms of (the ordinary locus of) the curve with some Lubin-Tate p -divisible group.

Theorem 2.9. *For every $\varpi \geq 2$, $g \geq 2$ such that $p > 4^{3g-3}$, the Lubin-Tate stack \mathcal{L}_g^ϖ is nonempty.*

Proof. Note that the case $\varpi = 1$ was discussed in [Mzk1] by studying the indigenous bundles on a totally degenerate curve. We shall take a similar approach here, only the situation is more delicate for a number of technical reasons. Thus, let S^{\log} be a log scheme whose underlying scheme is $\mathrm{Spec}(W(k)/p^{\varpi}W(k))$ (where k is an algebraically closed field of characteristic p). Let $X^{\log} \rightarrow S^{\log}$ be a totally degenerate (0-pointed) stable curve of genus g . Now we would like to apply the theory of §2 of Chapter II, together with Theorem 2.3. First we claim the following:

(*) For an appropriate choice of totally degenerate curve X^{\log} , there exists an indigenous bundle (P, ∇_P) on X^{\log} which defines a nilpotent, nondormant n -connection (where $n \stackrel{\mathrm{def}}{=} \varpi - 1$) on (P_0, ∇_{P_0}) .

Indeed, this is clear when $g = 2$: we simply take two copies of the 3-pointed curve of genus 0 equipped with the totally indigenous bundle discussed in the first Remark following Theorem 2.3, and glue them together along the marked points with corresponding radii. For arbitrary $g \geq 2$, one observes that since the fundamental group of the dual graph of a totally degenerate curve of genus 2 is the free group on two generators, it has finite quotients of arbitrary order; thus, the curve appearing in the example just constructed in the case $g = 2$ has a finite, étale covering by a totally degenerate curve of arbitrary genus $g \geq 2$, so by pulling back the indigenous bundle we constructed downstairs, we get an indigenous bundle on our X^{\log} of arbitrary genus $g \geq 2$ upstairs that has the desired properties.

Next, let us recall the discussion preceding Proposition 2.10 of Chapter II. In the notation of this discussion, it was shown in Proposition 2.10 that there exist complete noetherian local rings $R'_\xi \stackrel{\mathrm{def}}{=} \widehat{\mathcal{O}}_{N_S, \xi^p}$ and R_ξ with residue field k such that R'_ξ has the structure of a R_ξ -algebra whose structure morphism $R_\xi \rightarrow R'_\xi$ is a local, finite, flat, local complete intersection morphism of degree $\leq 4^{3g-3}$, and that R_ξ is equipped with a formally smooth morphism $\mathrm{Spec}(R_\xi) \rightarrow (\overline{\mathcal{M}}_g)_k$ that maps the closed point of $\mathrm{Spec}(R_\xi)$ to the point of $(\overline{\mathcal{M}}_g)_k$ corresponding to the totally degenerate curve of the preceding paragraph. Thus, by taking “sections” of the morphism $\mathrm{Spec}(R_\xi) \rightarrow (\overline{\mathcal{M}}_g)_k$, it is clear that we may construct a fine noetherian log scheme U^{\log} and a stable log-curve $Y_0^{\log} \rightarrow U_0^{\log}$ of genus g with the following properties:

- (1) U is a connected, $\mathbf{Z}/p^{n+1}\mathbf{Z}$ -flat S -scheme of finite type such that the classifying morphism $U_0 \rightarrow (\overline{\mathcal{M}}_g)_k$ (defined by $Y_0^{\log} \rightarrow U_0^{\log}$) is quasi-finite and flat, and every irreducible component of U_{red} is *generically étale* (here we use that $p > 4^{3g-3}!$) over $(\overline{\mathcal{M}}_g)_k$;
- (2) there exists a crystal (Q, ∇_Q) in \mathbf{P}^1 -bundles on $\mathrm{Crys}(Y_0^{\log}/U^{\log})$ that defines a nilpotent n -connection on

(Q_0, ∇_{Q_0}) (which we assume to be indigenous on Y_0^{\log}) whose p^{n+1} -curvature does not vanish identically on any fiber of $Y_0 \rightarrow U_0$;

- (3) there exists a morphism $S^{\log} \rightarrow U^{\log}$ (such that $S \rightarrow U$ is an S -morphism) with the property that the pull-back of (Q, ∇_Q) (respectively, $Y_0^{\log} \rightarrow U_0^{\log}$) to S^{\log} via this morphism is (P, ∇_P) (respectively, $X_0^{\log} \rightarrow S_0^{\log}$).

Note that in general, U_0 may not be reduced. Indeed, this is because U_0 is a sort of algebraization of the inverse image in $\mathrm{Spec}(R'_\xi)$ of a “section” of $\mathrm{Spec}(R'_\xi) \rightarrow (\overline{\mathcal{M}}_g)_k$, and we must use the fact that $R'_\xi \rightarrow R'_\xi$ is a local complete intersection morphism in order to obtain the $\mathbf{Z}/p^{n+1}\mathbf{Z}$ -flat lifting U of U_0 .

Next, we would like to consider the issue of whether or not the p^{n+1} -curvature of (Q, ∇_Q) has any zeroes. Let $u \in U(k)$ be a point over which the curve $Y_0 \rightarrow U_0$ is smooth. Let $Y_u \rightarrow \mathrm{Spec}(k)$ be the fiber of $Y_0 \rightarrow U_0$ over u . The p^{n+1} -curvature of (Q, ∇_Q) at u may be regarded as a morphism

$$\mathrm{Ad}(Q_0) \rightarrow (\Phi_{Y_u}^*)^{n+1} \omega_{Y_u/k}$$

Let $Z \subseteq Y_u$ be the zero locus of this morphism. Since (Q_0, ∇_{Q_0}) is *crys-stable*, and the p^{n+1} is not identically zero (cf. (2) above), it follows that Z is a divisor in Y_u and

$$\deg(Z) < p^{n+1} \cdot (2g - 2)$$

Moreover, it follows from the horizontality properties of the p^{n+1} -curvature that $Z \subseteq Y_u$ is the inverse image via $\Phi_{Y_u}^{n+1}$ of a divisor $Z' \subseteq Y_u$. Thus, $\deg(Z') < 2g - 2$. Let us call the \mathbf{P}^1 -bundle $Q_0 \rightarrow Y_0$ *even* (respectively, *odd*) if there exists (respectively, does not exist) a line bundle (étale locally on U_0) whose square is equal to τ_{Q_0/Y_0} . Since $Q_0 \rightarrow Y_0$ is indigenous and $2g - 2$ is even, it is easy to see that $Q_0 \rightarrow Y_0$ is even. On the other hand, since $(\Phi_{Y_u}^*)^{n+1} \omega_{Y_u/k}$ has even degree, if Z' had odd degree, it would follow that $Q_0|_{Y_u} \rightarrow Y_u$ is *odd*. Thus, we conclude that Z' has even degree. In particular, if $g = 2$, we thus obtain that $\deg(Z')$ is even and < 2 , hence zero, i.e., the p^{n+1} -curvature of (Q, ∇_Q) has no zeroes on Y_u .

For arbitrary $g \geq 2$, by using étale coverings as in the proof of (*) above, it follows that we can always assume that the data above has been constructed so that there exist $u \in U(k)$ for which the p^{n+1} -curvature of (Q, ∇_Q) has no zeroes on Y_u . Thus, in what follows we assume furthermore that:

- (4) there exists a $u \in U(k)$ such that Y_u is smooth and the p^{n+1} -curvature of (Q, ∇_Q) has no zeroes on Y_u .

Since the condition that the p^{n+1} -curvature of (Q, ∇_Q) have no zeroes on Y_u is an open condition on u , it follows that (by (1) above) we may assume that the u in (4) also has the property that $U_{\text{red}} \rightarrow (\overline{\mathcal{M}}_g)_k$ is *étale* at u . It is now a formal exercise in sorting through the definitions (i.e., sorting through the differences between the nilpotent n -connection point of view and the Lubin-Tate stack point of view) that there is an open neighborhood V of u in U_{red} such that the morphism $V \rightarrow (\overline{\mathcal{M}}_g)_k$ is *étale* and factors through $\overline{\mathcal{N}}_{g,r}^{\Pi,s}$. Thus, we get a morphism $V \rightarrow (\overline{\mathcal{N}}_{g,r}^{\Pi,s})_k$. Moreover, the fact that $(\overline{\mathcal{N}}_{g,r}^{\Pi,s})_k$ is smooth over k (by Proposition 2.7) implies that $(\overline{\mathcal{N}}_{g,r}^{\Pi,s})_k \rightarrow (\overline{\mathcal{M}}_g)_k$ is *étale* at the image of $V \rightarrow (\overline{\mathcal{N}}_{g,r}^{\Pi,s})_k$, as desired. This completes the proof. \circ

Remark. In fact, it follows from the above proof that for every $g \geq 2$, $\varpi \geq 2$, without any assumption on p , the VF-stack \mathcal{N}_g^{Π} of pure tone ϖ (where we omit the “bar” from \mathcal{N} since this VF-stack lies over \mathcal{M}_g) has a nonempty open substack that is quasi-finite over \mathcal{M}_g . That is to say, the assumption $p > 4^{3g-3}$ is used only to show that there exists a nonempty open substack of \mathcal{N}_g^{Π} that is, in fact, *étale* over \mathcal{M}_g .

§3. Anabelian Stacks

In this §, we study the “opposite phenomenon” (i.e., few zeroes in a period) to that of a VF-pattern of pure tone (i.e., lots of zeroes in a period). Moreover, we show the nonemptiness of certain “spiked VF-stacks” which, in some sense, constitute the most complicated of the various VF-stacks discussed in this book: that is to say, they represent a sort of mixture of pre-home stacks and Lubin-Tate stacks.

§3.1. Basic Definitions

Let Π be a VF-pattern of period ϖ . Then we make the following

Definition 3.1. We shall call Π *anabelian* if a period of Π (i.e., the image under Π of the set $\{0, \dots, \varpi - 1\}$) contains more than one positive level. We shall call Π *spiked* if the image of Π contains a positive level other than χ .

A typical example of an anabelian Π is the pre-home VF-pattern of period $\varpi \geq 2$. In this case, $\overline{\mathcal{N}}_{g,r}^{\Pi}$ is the fibered product of ϖ copies

$\overline{\mathcal{N}}_{g,r}^{\text{adm}}$ over $\overline{\mathcal{M}}_{g,r}$. Thus, it follows (by [Mzk1], Chapter II, Theorem 2.3, Corollary 2.15, and Corollary 3.8) that $\overline{\mathcal{N}}_{g,r}^{\Pi}$ is nonempty and quasi-finite and flat over $\overline{\mathcal{M}}_{g,r}$. Moreover, as we saw in Theorem 2.10 of Chapter III, $\overline{\mathcal{N}}_{g,r}^{\Pi}$ is affine.

Suppose that Π is spiked. Suppose that $\Pi(i-1)$ is a positive level other than χ . Let (P_i, ∇_{P_i}) be the i^{th} link bundle of some S -valued point of $\overline{\mathcal{N}}_{g,r}^{\Pi}$. Then the p -curvature locus of (P_i, ∇_{P_i}) is a $\Pi(i-1)$ -balanced divisor, hence, in particular, nonempty. We shall refer to the points of the Frobenius pull-back of the p -curvature locus of (P_i, ∇_{P_i}) as *spikes of (P_i, ∇_{P_i})* . In the remainder of this §, we would like to show the existence of spiked Π for which $\overline{\mathcal{N}}_{g,r}^{\Pi}$ is nonempty.

§3.2. Nondormant Bundles on the Projective Line Minus Three Points

Let k be an algebraically closed field of odd characteristic p . Let $S = \text{Spec}(k)$, and endow S with the trivial log structure. Let $X^{\log} \rightarrow S^{\log}$ be the 3-pointed smooth curve of genus 0. Thus, $X = \mathbf{P}_k^1$, and we think of the three marked points as 0, 1, and ∞ . Let $\rho = \{\rho_0, \rho_1, \rho_{\infty}\}$ be a set of three radii in \mathbf{F}_p . Let (P, ∇_P) be the torally indigenous bundle on X^{\log} of radii ρ . Let us refer to the quantity “twice the radius” as the *diameter*. Thus, (P, ∇_P) has diameters given by 2ρ .

Note that $P \rightarrow X$ is necessarily an *odd* \mathbf{P}^1 -bundle (i.e., it is the projectivization of a rank two vector bundle of odd degree). Indeed, this follows from the fact that the canonical height of the Hodge section $h : X \rightarrow P$ is $\frac{1}{2}$ an *odd* integer. Let $\mathcal{T} = \Phi_X^* \tau_{X^{\log}/S^{\log}}$; endow \mathcal{T} with its usual connection $\nabla_{\mathcal{T}}$ of p -curvature zero. Let us consider the p -curvature of (P, ∇_P) :

$$\mathcal{P} : \mathcal{T} \rightarrow \text{Ad}(P)$$

Note that because we are dealing with the case $g = 0$, $r = 3$, the p -curvature \mathcal{P} is *automatically nilpotent*. Suppose, however, that (P, ∇_P) is *not dormant*.

Let us consider the zero locus $V_{\mathcal{P}} \subseteq X$ of the p -curvature. Since \mathcal{T} has odd degree, and the \mathbf{P}^1 -bundle $P \rightarrow X$ is odd, it follows that $V_{\mathcal{P}}$ must be a divisor of *even degree* on X . Also, if we write $V_{\mathcal{P}}$ in the form $V(\mathcal{I})$, where $\mathcal{I} \subseteq \mathcal{O}_X$ is an ideal, then the ideal \mathcal{I} must be stabilized by the trivial logarithmic connection $\mathcal{O}_X \rightarrow \omega_{X/S}^{\log}$ on \mathcal{O}_X and have monodromy equal to $\pm 2\rho_i$ at the marked point labeled i (where $i \in \{0, 1, \infty\}$). Concretely, this means that any nonmarked point in the support of $V_{\mathcal{P}}$ has multiplicity (in $V_{\mathcal{P}}$) divisible by p , while the marked point labeled i has multiplicity congruent to $\pm 2\rho_i$ modulo p . Note that if $V_{\mathcal{P}}$ has degree $\geq p = \deg(\mathcal{T})$, then we obtain a contradiction to the

crys-stability of $P \rightarrow X$ (since the Hodge section $h : X \rightarrow P$ has negative canonical height), so $\deg(V_P) \leq p-1$. It thus follows that all the points in the support of V_P are marked points. Finally, note that if $\rho_i = 0$ (for some i), then it follows by direct computation of the p -curvature (see the discussion preceding Proposition 1.3 of Chapter II) that the marked point i is not contained in the support of V_P . We summarize these observations as follows:

Theorem 3.2. *Let (P, ∇_P) be a torally indigenous bundle of radii ρ on X^{\log} (the projective line with three marked points). Then (P, ∇_P) is automatically nilpotent, and its isomorphism class is uniquely determined by ρ . If it is not dormant, then the zero locus V_P of its p -curvature is a divisor of even degree $\leq p-1$ on X whose (set-theoretic) support is contained in the divisor consisting of those marked points with nonzero radii. Moreover, in this case, the degree of V_P at a marked point i is congruent modulo p to $\pm 2\rho_i$.*

Finally, if (P, ∇_P) is nondormant, then the divisor V_P is uniquely determined by the description in the preceding paragraph.

Proof. It remains only to prove the final uniqueness assertion. But this follows from the following elementary observation: If a, a', b, b', c, c' are nonnegative integers $\leq p-1$ such that $a \equiv \pm a'$, $b \equiv \pm b'$, $c \equiv \pm c'$ (modulo p), and $a+b+c$, $a'+b'+c'$ are both even and $\leq p-1$, then it follows that $a = a'$, $b = b'$, and $c = c'$. Indeed, since $a+b+c$ and $a'+b'+c'$ are both even, it follows that if it is not the case that $a = a'$, $b = b'$, $c = c'$, then (up to permuting “ a ,” “ b ,” and “ c ”) we may assume that $a = p-a'$, $b = p-b'$, and $c = c'$. But then $2(p-1) \geq (a+b+c) + (a'+b'+c') = p+p+2c \geq 2p$, so $p-1 \geq p$, which is absurd. This completes the proof of Theorem 3.2. \circ

Thus, together with Theorem 2.3, Theorem 3.2 implies the following:

Given a triple of radii in \mathbf{F}_p , it is possible to completely determine, just from looking at ρ , whether the corresponding (P, ∇_P) is dormant or not, and, if it is nondormant, what its V_P looks like.

Next, let us recall the discussion of *adjustability* of §1.6 of Chapter II. First, we suppose that the radius of (P, ∇_P) at the marked point i is equal to 1. We will say that (P, ∇_P) is *adjustable* (cf. Definition 1.13 of Chapter II) at the marked point i if, when one writes (P, ∇_P) as the projectivization of some $(\mathcal{E}, \nabla_{\mathcal{E}})$ with trivial determinant in a neighborhood of i , then there exists a local section s of \mathcal{E} such that

$$\nabla(s) \equiv s \text{ modulo } t^3$$

(where t is a local coordinate on X at i that vanishes at i , and ∇ is the result of applying $\nabla_{\mathcal{E}}$ in the direction “ $t\partial/\partial t$ ”). (It is easy to see that this is just an abridged version of Definition 1.13 of Chapter II.)

Now let us consider unique horizontal sub-line bundle $\mathcal{L} \subseteq \mathcal{E}$ (i.e., such that the inclusion $\mathcal{L} \hookrightarrow \mathcal{E}$ is locally split) whose existence follows from the fact that (P, ∇_P) is nilpotent. If the multiplicity of i in V_P is $p - 2$ (respectively, 2), then, when we restrict to the marked point i , the subspace $\mathcal{L}|_i \subseteq \mathcal{E}|_i$ is the eigenspace with eigenvalue 1 (respectively, -1) for the monodromy operator associated to $\nabla_{\mathcal{E}}$. Thus, we conclude that:

If $\rho_i = 1$ (for some marked point i), then the marked point i is *adjustable* (in the sense of Definition 1.13 of Chapter II) if and only if the multiplicity of i in V_P is $p - 2$ (i.e., not 2).

In fact, the main example of interest in the construction of §3.3 below is the case where the set of diameters is $\{0, 1, 2\}$. In this case, (P, ∇_P) is adjustable at the marked point α of diameter 2. Moreover, if one *adjusts* (as in §1.6 of Chapter II) this (P, ∇_P) at α , one obtains a \mathbf{P}^1 -bundle with connection (Q, ∇_Q) on the 2-pointed semi-stable log-curve Y^{\log} obtained from X^{\log} by letting α be unmarked. Then the Hodge section defines a locally split injection $\omega_{Y^{\log}}(-\alpha) \hookrightarrow \text{Ad}(Q)$, while the p -curvature defines a morphism

$$\mathcal{P}_Q : \text{Ad}(Q) \rightarrow (\Phi_Y^* \omega_{Y^{\log}})(-\beta)$$

(where β is the marked point of Y^{\log} of diameter 1). Note, moreover, that \mathcal{P}_Q is *surjective*. Indeed, this is clear away from α . On the other hand, if \mathcal{P}_Q were zero at α , then $\text{Ad}(Q)$ would admit a horizontal surjection onto a line bundle of degree $\leq -p - 1$ which would imply (since $\deg(\omega_{Y^{\log}}(-\alpha)) = -1$) that the Hodge section is horizontal. But this is absurd. This contraction completes the proof that \mathcal{P}_Q is surjective. In particular, it follows that *the p -curvature of (Q, ∇_Q) is nonzero at α .*

The following technical consequences of the above discussion will be useful in the sequel:

Lemma 3.3. *If $2\rho = \{0, 1, 2\}$, then V_P has multiplicity 0 (respectively, 1; $p - 2$) at the marked point whose diameter is 0 (respectively, 1; 2). In particular, in this case, (P, ∇_P) is adjustable at the marked point of diameter 2. Moreover, if α is the marked point of diameter 2, and Y^{\log} is the 2-pointed semi-stable log-curve obtained from X^{\log} by letting α be unmarked, then the p -curvature of the \mathbf{P}^1 -bundle with connection on Y^{\log} obtained by adjusting (P, ∇_P) at α is nonzero at α .*

Conversely, if $2\rho = \{0, d, 2\}$, and (P, ∇_P) has a marked point of diameter 2 at which it is adjustable, then $d = \pm 1$.

Lemma 3.4. *If $2\rho = \{0, \lambda, \lambda\}$, where λ is an odd positive integer $\leq \frac{1}{2}(p-1)$, then $V_{\mathcal{P}}$ has multiplicity 0 (respectively, λ) at the marked point(s) whose diameter is 0 (respectively, λ).*

Lemma 3.5. *If $2\rho = \{0, 0, \lambda\}$, where λ is a nonnegative even integer $\leq p-1$, then $V_{\mathcal{P}}$ has multiplicity 0 (respectively, λ) at the marked point(s) whose diameter is 0 (respectively, λ).*

§3.3. Explicit Construction of Spiked Data

Let k and $X^{\log} \rightarrow S^{\log}$ be as in the preceding subsection. Let Y^{\log} and Z^{\log} be copies of X^{\log} . Let W be the 5-pointed stable curve of genus zero constructed as follows (cf. the Pictorial Appendix): We glue X to Y by identifying ∞_X (i.e., the point “ ∞ ” on X) to 0_Y , and Y to Z by identifying ∞_Y to 0_Z . Thus, W has three irreducible components (X , Y , and Z), and 5 marked points (0_X , 1_X , 1_Y , 1_Z , and ∞_Z). In particular, W defines a point of $\overline{\mathcal{M}}_{0,5}(k)$. Let T^{\log} be the log scheme with $T = \text{Spec}(k)$, and the log structure obtained by pulling back to $\text{Spec}(k)$ the canonical log structure on $\overline{\mathcal{M}}_{0,5}$. Let

$$W^{\log} \rightarrow T^{\log}$$

be the log-curve obtained by pulling back the log structure on the tautological log-curve over $\overline{\mathcal{M}}_{0,5}^{\log}$.

Let (Q, ∇_Q) be the indigenous bundle on W^{\log} whose diameters at ∞_X , 0_Y , ∞_Y , and 0_Z are all 1. Thus, by Lemma 3.5, the multiplicity of $V_{\mathcal{P}|_X}$ (respectively, $V_{\mathcal{P}|_Z}$) at ∞_X (respectively, 0_Z) is $p-1$; by Lemma 3.4, the multiplicity of $V_{\mathcal{P}|_Y}$ at 0_Y and ∞_Y is 1 (cf. the Pictorial Appendix). It thus follows from the discussion preceding Proposition 3.3 of Chapter II that (Q, ∇_Q) is *mildly spiked of strength $2p$* . Let M be the trait (over $\text{Spec}(k)$) which is the completion at the generic point of the exceptional divisor of the blow-up of $(\overline{\mathcal{M}}_{0,5})_k$ at the point defined by W^{\log} . Endow M with the log structure pulled back from the canonical log structure on $\overline{\mathcal{M}}_{0,5}$, so as to obtain M^{\log} . Pulling back the tautological curve, we thus obtain a log-curve

$$W_M^{\log} \rightarrow M^{\log}$$

Moreover, by Proposition 3.10 of Chapter II, we see that (Q, ∇_Q) deforms uniquely to a mildly spiked indigenous bundle (of strength $2p$) (Q_M, ∇_{Q_M}) on W_M^{\log} .

Let η_M be the generic point of M . It is clear that the p -curvature locus of $Q_M|_{\eta_M}$ forms a divisor in $D_{\eta} \subseteq W_M^F|_{\eta_M}$ which is étale of degree 2 over η_M . Moreover, it follows from the fact that the two spikes that

appear in the generic fiber are genericizations of two distinct nodes in the special fiber that in fact, D_η splits over η_M . Thus, we obtain two sections $\delta, \epsilon : M \rightarrow W_M^F$ such that $(\delta \cup \epsilon)|_{\eta_M} = D_\eta$.

Clearly, δ and ϵ map the special point of M to the two nodes of $(W_M^{\log})^F$. Thus, by blowing up $(W_M^{\log})^F$, we may form a 7-pointed curve of genus zero

$$U_M^{\log} \rightarrow M^{\log}$$

such that $U_M^{\log}|_{\eta_M}$ is given by $(W_M^{\log})^F|_{\eta_M}$ equipped with δ and ϵ as its sixth and seventh marked points. Note that we have a blow-down morphism $U_M^{\log} \rightarrow (W_M^{\log})^F$, given by forgetting the marked points δ and ϵ . Let U^{\log} be the special fiber of U_M^{\log} . Then U^{\log} is obtained from $(W^{\log})^F$ as follows (cf. the Pictorial Appendix): First let us identify W^{\log} with $(W^{\log})^F$, since W^{\log} is clearly defined over \mathbf{F}_p . Let D^{\log} and E^{\log} be copies of X^{\log} . Then U^{\log} is given by gluing X^{\log} at ∞_X to D^{\log} at 0_D ; D^{\log} at ∞_D to Y^{\log} at 0_Y ; Y^{\log} at ∞_Y to E^{\log} at 0_E ; and E^{\log} at ∞_E to Z^{\log} at 0_Z . Here, the image of the special point of M under δ (respectively, ϵ) is 1_D (respectively, 1_E). Moreover, the morphism $U^{\log} \rightarrow W^{\log} = (W^{\log})^F$ is given by contracting D^{\log} and E^{\log} to a point.

Now we would like to construct a torally indigenous bundle (R, ∇_R) on U^{\log} as follows (cf. the Pictorial Appendix): (R, ∇_R) is to have

- (1) diameter 0 at $0_X, 1_X, \infty_X \sim 0_D, 1_Y, \infty_E \sim 0_Z, 1_Z$, and ∞_Z ;
- (2) diameter 1 at $\infty_D \sim 0_Y$ and $\infty_Y \sim 0_E$; and
- (3) diameter 2 at 1_D and 1_E .

This determines (R, ∇_R) . As observed in Lemma 3.3, (R, ∇_R) is adjustable at 1_D and 1_E . Thus, by Theorem 1.16 of Chapter II, after replacing M by some trait L (which is finite and flat over M , and equipped with the log structure defined by its special point), we obtain that (R, ∇_R) deforms to a nilpotent torally indigenous bundle (R_L, ∇_{R_L}) over U_L^{\log} (i.e., the pull-back to L^{\log} of U_M^{\log}), which is adjustable at δ_L and ϵ_L . Moreover, if $p \geq 5 > 2^2$, then we may take L to be generically étale over M . (Note that here the reader might wonder, relative to the statement of Theorem 1.16 of Chapter II, why $p > 2^2$, rather than $p > 2^4$ is sufficient. The reason is that relative to this issue of generic étaleness, one can forget about the components “ X ” and “ Y ” of U ; this yields the sharper bound $p > 2^2$.)

Remark. Because of the relatively simple structure of the curve U^{\log} (in particular, the fact that it remains semi-stable even when we regard the sixth and seventh marked points as unmarked), it is easy to see

that the proof of Theorem 1.16 in Chapter II can be considerably simplified (i.e., one no longer needs Lemma 1.14 (of Chapter II), and the proof of Lemma 1.15 (Chapter II) becomes trivial) in the special case of the curve U_M^{\log} . We leave the details to the reader.

Let $V_L^{\log} \rightarrow L^{\log}$ be the semi-stable 5-pointed curve of genus zero obtained from U_L^{\log} by forgetting δ and ϵ . Let V^{\log} be its special fiber. Let (P, ∇_P) be the bundle on V_L^{\log} obtained by adjusting (R_L, ∇_{R_L}) at δ_L and ϵ_L . Thus, the restriction $(P, \nabla_P)_{\eta_L}$ of (P, ∇_P) to the generic fiber of V_L^{\log} is crys-stable of level $\frac{1}{2}$. Now we claim that $(P, \nabla_P)_{\eta_L}$ is *admissible*. Indeed, if it were not admissible, it would follow that (after possibly replacing L by a finite flat trait over L) its p -curvature vanishes at some divisor $B_{\eta_L} \subseteq (V_L)_{\eta_L}$ of degree p whose support is concentrated at a single point of $(V_L)_{\eta_L}$. Let $B \subseteq V_L$ be the schematic closure of B_{η_L} in V_L . It would thus follow that the p -curvature of $(P, \nabla_P)|_V$ must vanish on $B|_V$, i.e., on some closed subscheme of length p in V which is concentrated at a single point. But this is absurd, as one sees from the construction of (R, ∇_R) (and the description of the multiplicities of the zero locus of the p -curvature in Lemmas 3.3 and 3.4) – cf. also the Pictorial Appendix. This completes the proof of the claim that $(P, \nabla_P)_{\eta_L}$ is admissible.

It thus follows from everything that we have done that $(Q_L, \nabla_{Q_L})_{\eta_L}$ and $(P, \nabla_P)_{\eta_L}$ form an η_L -valued point of $\overline{\mathcal{N}}_{0,5}^{\Pi,s}$, where Π is the VF-pattern of period 2 such that $\Pi(0) = \frac{3}{2}$, and $\Pi(1) = \frac{1}{2}$. In particular, we see that we have proven the following:

Lemma 3.6. *Let $p \geq 5$, $g = 0$, $r = 5$, $\varpi = 2$. Let Π be the VF-pattern of period ϖ such that $\Pi(0) = \chi = \frac{3}{2}$, $\Pi(1) = \frac{1}{2}$. Then the shifted VF-stack $\overline{\mathcal{N}}_{g,r}^{\Pi,s}$ contains a nonempty open substack U such that U_{red} is étale over $\overline{\mathcal{M}}_{g,r}$.*

More generally, we have the following result, which is the main result of this §:

Theorem 3.7. *Consider the following two statements: (i) there exists a spiked VF-pattern Π of period 2 for which $\overline{\mathcal{N}}_{g,r}^{\Pi,s}$ is nonempty; (ii) $2g - 2 + r \geq 3$. Then (i) implies (ii); moreover, if $p \geq 5$, then (ii) implies (i).*

Finally, if $p \geq 5$ and $2g - 2 + r \geq 3$, then we can choose the VF-pattern so that $\overline{\mathcal{N}}_{g,r}^{\Pi,s}$ contains a nonempty open substack U such that U_{red} is étale over $\overline{\mathcal{M}}_{g,r}$.

Proof. The necessity of the condition $2g - 2 + r \geq 3$ follows from the fact that if (P, ∇_P) is a mildly spiked indigenous bundle on a smooth curve X^{\log} with $2g - 2 + r \leq 2$, then the dual p -curvature $\text{Ad}(P) \rightarrow \Phi_{X^{\log}}^* \omega_{X/S}^{\log}$ must

(by a parity argument) have a zero locus of even degree divisible by p ; but if this degree is nonzero, then we obtain a nonzero horizontal morphism of $\mathrm{Ad}(P)$ onto a line bundle of degree zero, so the crystal stability of (P, ∇_P) is violated.


As for sufficiency (when $p \geq 5$), the case $g = 0$; $r = 5$ was handled in Lemma 3.6, so we may simply do the following:

- (1) take a disjoint union Y^{\log} of totally degenerate stable pointed curves equipped with 0^{th} and 1^{st} link bundles given by the unique indigenous bundle which is admissible (i.e., all the radii are zero);
- (2) glue the data of (1) to the example of Lemma 3.6 at various marked points of the log-curve of that example.


(Here, one chooses Y^{\log} appropriately so that the curve resulting from the gluing has the desired g and r . Of course, for some (g, r) , it may be necessary to take Y^{\log} to be the “empty set curve” and/or to glue marked points of the log-curve of the example of Lemma 3.6 to other marked points of the log-curve of the example of Lemma 3.6; we leave such combinatorial details to the reader.) It is immediate that by deforming the resulting curve/bundles to a smooth curve with appropriate bundles, one obtains the data whose existence is claimed in the Theorem. \circ

Pictorial Appendix


First we give a picture of the curve W :




Now we give an illustration of the multiplicities of the zero locus of the p -curvature for the bundle (Q, ∇_Q) on W^{\log} :



Next, we give a picture of the curve U :



Finally, we give an illustration of the multiplicities of the zero locus of the p -curvature for the bundle (R, ∇_R) on U^{\log} :



Chapter V: Combinatorialization at Infinity of the Stack of Nilcurves

§0. Introduction

In this Chapter, we continue our study of $\overline{\mathcal{N}}_{g,r}^\rho$. Our main result is Theorem 1.1, stated at the end of §1. Unfortunately, because of the technically complicated nature of the result, in order even to state Theorem 1.1 precisely, it is necessary to introduce a large number of technical terms, which we do in §1. In the present §, we seek, by contrast, to give a general overview of what is going on, eschewing as far as is possible the use of precise technical terms as in §1.

In some sense, the most fundamental object in the theory of [Mzk1] and the present book is the stack $\overline{\mathcal{N}}_{g,r}^\rho$ of r -pointed stable curves (in characteristic p) of genus g equipped with a nilpotent torally indigenous bundle of radii ρ . In the following discussion, it will be convenient to refer to stable curves equipped with a torally indigenous bundle (respectively, nilpotent torally indigenous bundle) as *indigenized* (respectively, *nilindigenized*) curves. In fact, we shall even refer to nilindigenized curves simply as *nilcurves* for short. A nilcurve whose underlying curve is a 3-pointed curve of genus zero will be called an *atom*. A nilcurve whose underlying curve is maximally degenerate will be called a *molecule*.

Of course, the structure of $\overline{\mathcal{N}}_{g,r}$ has already been studied in [Mzk1], in some detail, especially in Chapter II of [Mzk1]. Perhaps the most fundamental fact concerning $\overline{\mathcal{N}}_{g,r}$ (cf. Introduction, Theorem 0.1; [Mzk1], Chapter II, Theorem 2.3) is that it is *finite and flat of degree p^{3g-3+r} over $(\overline{\mathcal{M}}_{g,r})_{\mathbb{F}_p}$* . Moreover, in Chapters II and IV of the present book, one finds a complete description of precisely when an atom is *dormant* (i.e., its p -curvature is identically zero) (Chapter IV, Theorem

2.3), and if it is nondormant, precisely what the *vanishing locus of its p -curvature* looks like (Chapter IV, Theorem 3.2). That is to say,

The analysis of the structure of atoms has already been completed in Chapters II and IV.

In general, when studying objects on smooth hyperbolic curves, it is often useful to analyze what happens to such objects when they degenerate, i.e., when the underlying curves degenerate, especially when they degenerate to maximally degenerate curves. In the case of nilcurves, one is especially fortunate in that $\overline{\mathcal{N}}_{g,r}^\rho$ is finite and flat over $(\overline{\mathcal{M}}_{g,r})_{\mathbb{F}_p}$, so

In order to understand the generic points of $\overline{\mathcal{N}}_{g,r}^\rho$, it suffices to understand molecules, together with the issue of what sort of generic smooth nilcurve a given molecule deforms to.

Since molecules are just formed by concatenating atoms in a fashion compatible with their radii, the problem of analyzing all molecules has essentially already been solved in Chapters II and IV.

What is somewhat less trivial, however, is understanding how a given molecule deforms. In certain special cases, this problem was solved in [Mzk1], as well as in Chapter II, §2 and 3 of the present book. However, up till now, only the best behaved cases were dealt with, and, moreover, they were dealt with essentially from the point of view of proving the *existence* of smooth nilcurves of a given type, not from the point of view of accounting for *all* possible molecules that a given type of smooth nilcurve might degenerate to.

The purpose of the present Chapter, thus, is to give a complete answer (Theorem 1.1) to the question of what types of generic smooth nilcurve a given molecule can deform to. Here, by “types,” we mean whether the nilcurve is dormant or not, and, if it is nondormant, the degree of its spiked locus. From this result, we obtain as a corollary an entirely combinatorial algorithm for computing the generic degree (over $(\overline{\mathcal{M}}_{g,r})_{\mathbb{F}_p}$) of the locus of smooth nilcurves of a given type (cf. Corollary 1.3). We refer to §1 for precise statements of Theorem 1.1 and Corollary 1.3; and we refer to the Remark immediately following Theorem 1.1 for a more detailed description of the contents of this Chapter.

Finally, in keeping with the theme (present in the Introductions of [Mzk1] and the present book) of discussing analogies with complex Teichmüller theory, we note that

In some sense the notion of an “atom” discussed here is analogous to the notion of “pants” in complex Teichmüller theory – i.e.,

both form the most fundamental geometric building blocks of a hyperbolic curve.

Thus, there is a certain analogy between the theory of the present Chapter and the theory over the complex numbers that analyzes a hyperbolic Riemann surface in terms of how it is built up out of pants. Of course, in the complex case, the degeneration process is more transparent because “everything is ordinary” (i.e., acts like the case of a molecule made up of atoms all of whose radii are zero), and so one need not worry about the complications that arise from atoms with nonzero radii. Perhaps, however, there is a certain analogy between the Lubin-Tate portions of the geometries that can arise in the p -adic case (cf. Chapter VIII) and the case of hyperbolic Riemann surfaces with boundaries of nonzero radius in the complex case. If this analogy is to be believed, then perhaps it is appropriate to regard the order $\lambda \in \mathbf{Z}_{>0}$ (see Chapter VIII, §2.5) of a Lubin-Tate geometry as analogous to the positive real number which is the radius of a boundary (i.e., of a “pants” that makes up the Riemann surface in question). This analogy makes all the more sense if one recalls that $\mathbf{Z}_{>0}$ and $\mathbf{R}_{>0}$ both form the positive portions of the value groups of the respective local fields in question (cf. the discussion in the Introduction to this book, §4.6).

§1. Statement of Main Results

In this §, we state the main results of this Chapter. Unfortunately, in order to state these results precisely, it will be necessary to introduce a number of technical definitions, as well as to review the *structure theory of atoms*, which is implicit in Chapter IV. Throughout, p will be an *odd prime*.

We begin by reviewing the structure theory of atoms. Let S^{\log} be a fine noetherian log scheme, where $S = \operatorname{Spec}(k)$, and k is an algebraically closed field of characteristic p . Let $Z^{\log} \rightarrow S^{\log}$ be the 3-pointed (smooth) log-curve of genus zero. We shall label its marked points by the symbols α , β , and γ . Next, let $\rho = \{\rho_\alpha, \rho_\beta, \rho_\gamma\}$ be an (ordered) set of *radii* in k , and let (R, ∇_R) be a *totally indigenous bundle* (see Chapter I, §4.1, for definitions) on Z^{\log} of radii ρ . (Strictly speaking, the radii are elements of $k/\{\pm 1\}$ (i.e., equivalence classes of elements $\lambda \in k$, in which λ and $-\lambda$ are identified), as opposed to k , but for simplicity, we shall abuse notation and act as if they are elements of k .) The quantity 2 times a radius will be referred to as a *diameter*. Then it follows from the theory of Chapters I (§4), II, and IV that

- (1) The isomorphism class of (R, ∇_R) is uniquely determined by its set ρ of three radii. On the other hand,

given any set of three radii ρ in k , there always exists a torally indigenous bundle with those radii (cf. Chapter I, Theorem 4.4).

- (2) (R, ∇_R) is *nilpotent* (i.e., its p -curvature is nilpotent) if and only if its radii are in \mathbb{F}_p (Chapter II, §1.2). In this Chapter, we shall refer to Z^{\log} equipped with a nilpotent torally indigenous bundle as an *atom*. Thus, an atom may be specified by specifying three radii (or, equivalently, three diameters) in \mathbb{F}_p (strictly speaking, $\mathbb{F}_p/\{\pm 1\}$). Often, when all the radii are *nonzero*, we shall find it convenient to deal with three fixed nonnegative integers a , b , and c satisfying: (i) $a, b, c < p$; (ii) the ordered set $\{a, b, c\} \pmod{p}$ is equal (up to ± 1) to the set of diameters 2ρ ; (iii) $a + b + c$ is odd and $< 2p$. It is easy to see that for any set of radii in \mathbb{F}_p , such a triple $\{a, b, c\}$ always exists.
- (3) A nilpotent (R, ∇_R) is *dormant* (i.e., its p -curvature is identically zero) if and only if the following condition holds: all the radii are nonzero, and there exist three cyclic permutations σ_a , σ_b , and σ_c (in the automorphism group of some finite set) of orders a , b , and c , respectively, such that $\sigma_a \cdot \sigma_b = \sigma_c$. (See Chapter IV, Theorem 2.3, for more details.) As was pointed out to the author by A. Tamagawa, a trivial elementary argument shows that this condition is equivalent to the condition: $a + b > c$; $a + c > b$; $b + c > a$.
- (4) If (R, ∇_R) is nilpotent, but nondormant, then its p -curvature does not vanish, except possibly at the marked points, and, moreover, the vanishing orders v_α , v_β , v_γ (necessarily nonnegative integers!) of the p -curvature at the three marked points are uniquely determined by the following conditions: (i) $v_\alpha + v_\beta + v_\gamma$ is even and $< p$; (ii) the ordered set $\{v_\alpha, v_\beta, v_\gamma\} \pmod{p}$ is equal (up to sign) to the ordered set of diameters 2ρ . (See Chapter IV, Theorem 3.2, for more details.)

What is important about the above properties is that they completely reduce the issue of labeling atoms, as well as the subtle arithmetic issues of whether a given atom is dormant, or, if it is nondormant, what the orders of vanishing of its p -curvature are, to *entirely combinatorial terms*.

Now let us shift gears, and suppose that $X^{\log} \rightarrow S^{\log}$ is a *maximally degenerate* r -pointed stable log-curve of genus g , equipped with a nilpotent torally indigenous bundle (P, ∇_P) of radii ρ (in \mathbb{F}_p). We shall call

such a pair $(X^{\log}, (P, \nabla_P))$ a *molecule*. If we restrict (P, ∇_P) to some connected component Z^{\log} of the normalization of X^{\log} (where the marked points of Z^{\log} are the points mapping to nodes or marked points of X^{\log}), then we obtain a nilpotent torally indigenous bundle (R, ∇_R) on Z^{\log} , i.e., an *atom*. On the other hand, if we give on *each* connected component Z^{\log} of the normalization of X^{\log} a nilpotent torally indigenous bundle (R, ∇_R) in such a way that the radii at marked points that glue together to become nodes of X^{\log} coincide (as elements of $\mathbb{F}_p/\{\pm 1\}$), then these atoms glue together uniquely to form a molecule. That is to say, *the study of molecules can be reduced entirely to the study of atoms (plus some more combinatorics)*.

Let $M \stackrel{\text{def}}{=} \{X^{\log}, (P, \nabla_P)\}$ be a *nondormant molecule* (i.e., it has at least one nondormant atom) of radii ρ . We shall refer to as a *preplot* Π for M a numbering (via the digits $1, 2, 3, \dots$) of some subset ν_Π of the nodes of M such that every dormant atom of M has at least one node in ν_Π . We shall refer to the nodes of ν_Π as the (Π) -*active nodes* of M , and we shall write $|\nu_\Pi|$ for the number of active nodes of M . Given a preplot Π , we can define a log-curve X_i^{\log} (where $i = 1, \dots, |\nu_\Pi|$) as follows: X_i^{\log} is obtained by gluing together the connected components of the normalization of X^{\log} that support nodes numbered $\leq i$ at precisely those nodes numbered $\leq i$. By restricting the torally indigenous bundle (arising from M) on X^{\log} to each of the X_i^{\log} , we thus obtain various *multi-molecules* M_i (whose underlying log-curves are the X_i^{\log}). (Here, by “multi,” we mean that the X_i^{\log} might not be connected; thus, each M_i is a disjoint union of “molecules” (as defined above).)

We shall call a preplot Π a *plot* for M if

- (1) No M_i (for $i = 1, \dots, |\nu_\Pi|$) has a dormant connected component (i.e., every connected component of every M_i has at least one nondormant atom).
- (2) M_{i+1} contains a dormant atom that is not in M_i (for $i = 1, \dots, |\nu_\Pi| - 1$).

Note that if Π is a plot, then $|\nu_\Pi| = n_{\text{dor}}$ (where n_{dor} is the number of dormant atoms in M). Roughly speaking,

The ordering defined by a plot describes the order in which the nodes are to be deformed.

The reason that we only need this ordering on (what is, in general) a proper subset (i.e., ν_Π) of the set of all nodes is that only at *those* nodes does the deformation theory depend in an essential way on the order in which the nodes are deformed. That is to say, it is *precisely by respecting this ordering that we are able to give a precise, quantitative description of how the p-curvature deforms as we deform a nilcurve*.

Before continuing, we observe that:

For every nondormant molecule M , there always exists a plot Π .

Indeed, suppose that the first i nodes have been chosen for some $i < n_{\text{dor}}$. Then since $i < n_{\text{dor}}$, there exist dormant atoms in M that are not contained in M_i . Moreover, since M is *connected* and *nondormant*, there exists *at least one dormant atom* in M which is not contained in any M_i , but which touches either a nondormant atom or an atom in M_i (at some node ν). We take ν for our $(i+1)^{\text{st}}$ node. By induction, this completes the proof of existence of a plot Π .

We shall refer to a formal completion of a one-dimensional k -smooth subscheme of X at a marked point or node of X as a *niche* of X^{\log} (or M). Thus, a niche is either a completion of X at a marked point, or a completion of a branch of a node of X at that node. A *sign* σ at a niche of M is a choice of a symbol “+,” “−,” or “0,” to be regarded as associated to that niche. If $\lambda \in \mathbb{F}_p/\{\pm 1\}$ and σ is a sign, then $\sigma(\lambda)$ is defined to be the unique nonnegative integer $\leq p-1$ which is even (respectively, odd; zero) if σ is “+” (respectively, “−”; “0”). A *product of signs* $\sigma_1 \cdot \sigma_2$ is the sign of any product of real numbers λ_1 and λ_2 of signs σ_1 and σ_2 , respectively.

Let us assume that M is equipped with a plot Π . Then for each $i = 1, \dots, n_{\text{dor}}$, we would like to define (inductively) the notion of a *scenario* for M_i as follows: First of all, a scenario for M_i consists precisely of a collection of choices of sign at each of the niches of M_i that satisfy certain conditions. The conditions are defined inductively as follows:

- (1) For every niche on a *nondormant* atom of M_i , we assign the sign “+” (respectively, “−”; 0) if the p -curvature of that atom has a zero of even nonzero order (respectively, odd order; zero order) at that niche.
- (2) Suppose that a scenario Σ_i has been given for M_i . Note that there is always precisely one dormant atom in M_{i+1} that does not belong to M_i . Let us denote this atom by N_i . Note that M_{i+1} has precisely one node that sits on N_i . We shall refer to the two niches associated to this node as *critical*. Of the two critical niches, we shall call the niche on N_i (respectively, not on N_i) the *outer* (respectively, *inner*) *critical niche*. Now the sign of the inner critical niche has already been determined (and is necessarily nonzero); thus, we stipulate that Σ_{i+1} assign to the outer critical niche the sign opposite to that of the inner critical niche. Moreover, we stipulate that the signs at the two noncritical niches of N_i be nonzero and such that the product of the signs of the three niches of N_i be “−.”

This completes the definition of the notion of a “scenario for M_i ,” where $i = 1, \dots, n_{\text{dor}}$. Then a *scenario for M* is an assignment of signs to each of the niches of M whose restriction to $M_{n_{\text{dor}}}$ constitutes a scenario for $M_{n_{\text{dor}}}$ and whose restriction to every nondormant atom of M is as described in condition (1) above. Note that there are precisely $2^{n_{\text{dor}}}$ possible scenarios for M . Finally, if Σ is a scenario for M , then we shall write

$$\Sigma(\rho)$$

for the sum, over all marked points x of M , of the quantities $\sigma_x(2\rho_x)$, where σ_x is the sign assigned by Σ to the niche at x , and ρ_x is the radius of M at x .

Next, let us assume that our nondormant molecule M is equipped not only with a *plot* Π , but also with a *scenario* Σ . Then we would like to classify the nodes of M into various types, as follows: Let ν be a node of M .

- (1) If Σ assigns to both of the niches of ν the sign 0, then we shall call ν *classical ordinary*.
- (2) If the node ν belongs to the set ν_{Π} , then we shall call ν *grafted*.
- (3) If ν is not grafted, and Σ assigns opposite nonzero signs to the two niches of ν , then we shall call ν *philial*.
- (4) If Σ assigns the same nonzero sign to the two niches of ν , then we shall call ν *aphilial*.

This completes our classification of the nodes of M . It is easy to see that every node of M belongs to precisely one of these classes. We refer to the Pictorial Appendix for a graphic summary of these various types of nodes. Note, moreover, that whether a node is classical ordinary or not does not depend on Π or Σ , and whether a node is grafted or not may depend on Π , but is independent of Σ .

Let us write $n_{\text{cl}}(\Sigma)$ (respectively, $n_{\text{grf}}(\Sigma)$; $n_{\text{phl}}(\Sigma)$; $n_{\text{aph}}(\Sigma)$) for the number of classical ordinary (respectively, grafted; philial; aphilial) nodes of M . Also, let us write n_{tor} for the number of nodes of M at which the radius is nonzero. Thus, $n_{\text{tor}} = 3g - 3 + r - n_{\text{cl}}(\Sigma)$. Also, we have $n_{\text{grf}}(\Sigma) = n_{\text{dor}}$. In particular, $n_{\text{grf}}(\Sigma)$ and $n_{\text{cl}}(\Sigma)$ are *independent of the choice of scenario Σ* . In general, however, when $n_{\text{dor}} > 0$, $n_{\text{phl}}(\Sigma)$ and $n_{\text{aph}}(\Sigma)$ *depend heavily on the choice of scenario Σ* .

Finally, we need one more definition: If U is a connected noetherian scheme of dimension zero, then we shall refer to as the *padding*

degree of U the length of the artinian local ring $A \stackrel{\text{def}}{=} \Gamma(U, \mathcal{O}_U)$. If, moreover, the length of A is equal to 2^d , where d is the dimension of the Zariski cotangent space of A , then we shall call U *taut*.

We are now in a position to state the Main Theorem of this Chapter.

Theorem 1.1. *Let p be an odd prime. Let M be a molecule of radii ρ over an algebraically closed field k of characteristic p . Let \mathcal{N} be the completion of $(\overline{\mathcal{N}}_{g,r}^\rho)_k$ at the point defined by M . Let \mathcal{M} be the completion of $(\overline{\mathcal{M}}_{g,r})_k$ at the point defined by the underlying curve of M ; let $\overline{\eta}$ be the strict henselization of the generic point of \mathcal{M} . Thus, we have a natural morphism $\mathcal{N} \rightarrow \mathcal{M}$, which (by Proposition 1.8 of Chapter II, §1.3) is known to be finite and flat of degree equal to $2^{n_{\text{tor}}}$. We shall write $\mathcal{N}_{\overline{\eta}}$ for the schematic fiber of \mathcal{N} over the “geometric point” $\overline{\eta} \rightarrow \mathcal{M}$ of \mathcal{M} . Moreover:*

- (1) *Suppose that M is dormant. Then \mathcal{N}_{red} maps isomorphically to \mathcal{M} , and, in particular, is formally smooth over k of dimension $3g - 3 + r$. Moreover, $\mathcal{N}_{\overline{\eta}}$ is taut, of padding degree 2^{3g-3+r} . Finally, the generic nilcurve represented by $(\mathcal{N}_{\overline{\eta}})_{\text{red}}$ is dormant.*
- (2) *Suppose that M is nondormant. Fix a plot Π for M . Then for each of the $2^{n_{\text{dor}}}$ scenarios Σ of (M, Π) , there exists an open substack $\mathcal{N}_\Sigma \subseteq \mathcal{N}_{\overline{\eta}}$ such that: (i) if $\Sigma \neq \Sigma'$, then $\mathcal{N}_\Sigma \cap \mathcal{N}_{\Sigma'} = \emptyset$, and, moreover, every point of $\mathcal{N}_{\overline{\eta}}$ is contained in some \mathcal{N}_Σ ; (ii) every residue field of \mathcal{N}_Σ is separable over (hence equal to) $k(\overline{\eta})$, and, moreover, $(\mathcal{N}_\Sigma)_{\text{red}}$ is of degree $2^{n_{\text{aph}}(\Sigma)}$ over $\overline{\eta}$; (iii) each point of \mathcal{N}_Σ is taut, of padding degree $2^{n_{\text{phl}}(\Sigma)}$; (iv) the smooth nilcurve represented by any point of $(\mathcal{N}_\Sigma)_{\text{red}}$ is mildly spiked of strength $\Sigma(\rho) + p \cdot n_{\text{phl}}(\Sigma)$, and has toral-ordinary rank (see Definition 2.2) $3g - 3 + r - n_{\text{phl}}(\Sigma)$.*

In other words, the arithmetic-geometry issue of determining what sorts of generic smooth nilcurves degenerate to the given molecule M is reduced (by means of the terminology introduced above) to an entirely combinatorial issue.

In particular, we obtain the following consequence: Let d be either a nonnegative integer or the symbol ∞ . If $\overline{\mathcal{N}}_{g,r}^\rho[d] \subseteq \overline{\mathcal{N}}_{g,r}^\rho$ is, respectively, the substack of mildly spiked torally indigenous bundles of strength d when $d < \infty$, or, the substack of dormant torally indigenous bundles when $d = \infty$, then (one knows from the theory of Chapter II that) $\overline{\mathcal{N}}_{g,r}^\rho[d]$ is (either empty or) k -smooth of dimension $3g - 3 + r$, and the issue of computing the generic degree of $\overline{\mathcal{N}}_{g,r}^\rho[d]$ over $(\overline{\mathcal{M}}_{g,r})_{\mathbb{F}_p}$ is reduced (by (1) and (2) above) to an entirely combinatorial problem.

Remark. The proof of Theorem 1.1 is given by augmenting the deformation theory of Chapter II – which is sufficient only to handle the

dormant case, as well as the cases of *classical ordinary* and *phial* nodes – by first studying deformations at *aphilial nodes* (§2.1) via degeneration theory, and then studying deformations at *grafted nodes* (§2.2 and 2.3) by extending the deformation theory of Chapter II, §3. In §2.4, we use some subtle but elementary commutative algebra to show that what we have done so far is sufficient to prove Theorem 1.1. In §3.1, we apply the theory of aphilial nodes (§2) to compute the local structure of $\overline{\mathcal{N}}_{1,1}$ at infinity (Corollary 3.3). In §3.2, we apply Theorem 1.1 to compute various invariants of $\overline{\mathcal{N}}_{g,r}$ explicitly in certain special cases. The results of §3.2 are summarized below in Corollary 1.3. Finally, in a Pictorial Appendix, we give a graphic description of the various types of deformation phenomena that one must analyze in order to prove Theorem 1.1.

Corollary 1.2. $(\overline{\mathcal{N}}_{g,r})_{\text{red}}$ is generically étale over $(\overline{\mathcal{M}}_{g,r})_{\mathbb{F}_p}$.

Proof. This follows formally from Theorem 1.1. \square

Remark. Thus, Corollary 1.2 implies in particular that *every admissible nilcurve can be deformed to an ordinary nilcurve*. Note that Corollary 1.2 had previously only been obtained for large p (Chapter II, Proposition 1.8).

Corollary 1.3. Let us denote by $n_{g,r}^{\text{ord}}$ the degree of $\overline{\mathcal{N}}_{g,r}^{\text{ord}}$ over $(\overline{\mathcal{M}}_{g,r})_{\mathbb{F}_p}$. (When it is necessary to specify p as well, we shall write $n_{g,r,p}^{\text{ord}}$.) Then the following hold:

- (1) Let $r \geq 3$ be an integer. Let $T_{0,r} \subseteq \mathbb{Z}^{r-3}$ be the set of (a_1, \dots, a_{r-3}) such that all the a_i are even and ≥ 0 ; and, moreover, for $i = 1, \dots, r-4$, we have $a_i + a_{i+1} \leq p-1$. If $\alpha = (a_1, \dots, a_{r-3}) \in T_{0,r}$, let N_α be the number of nonzero a_i 's. Then

$$n_{0,r}^{\text{ord}} = \int_{T_{0,r}} 2^{N_\alpha}$$

Let $\mathcal{T}_{0,r} \subseteq \mathbb{R}^{r-3}$ be the set of (a_1, \dots, a_{r-3}) such that all the a_i are ≥ 0 , and, moreover, for $i = 1, \dots, r-4$, we have $a_i + a_{i+1} \leq 1$. Then

$$\lim_{p \rightarrow \infty} \left(\frac{n_{0,r,p}^{\text{ord}}}{p^{r-3}} \right) = \int_{\mathcal{T}_{0,r}} 1$$

- (2) Let $r \geq 1$ be an integer. Let $T_{1,r} \subseteq \mathbb{Z}^r$ be the set of (a_1, \dots, a_r) such that all the a_i are ≥ 0 ; $a_1 + a_r$ is even and $\leq p-1$; and,

moreover, for $i = 1, \dots, r-1$, the quantity $a_i + a_{i+1}$ is even $\leq p-1$. If $\alpha = (a_1, \dots, a_r) \in T_{1,r}$, let N_α be the number of nonzero a_i 's. Then

$$n_{1,r}^{\text{ord}} = \int_{T_{1,r}} 2^{N_\alpha}$$

Let $\mathcal{T}_{1,r} \subseteq \mathbf{R}^r$ be the set of (a_1, \dots, a_r) such that all the a_i are ≥ 0 ; $a_1 + a_r \leq 1$; and, moreover, for $i = 1, \dots, r-1$, we have $a_i + a_{i+1} \leq 1$. Then

$$\lim_{p \rightarrow \infty} \left(\frac{n_{1,r,p}^{\text{ord}}}{p^r} \right) = \int_{\mathcal{T}_{1,r}} 2$$

(3) We have the following formulas:

$$\begin{aligned} n_{0,4}^{\text{ord}} &= p \\ n_{1,1}^{\text{ord}} &= p \\ n_{0,5}^{\text{ord}} &= \frac{1}{2}(p^2 + 1) \\ n_{1,2}^{\text{ord}} &= p^2 \\ n_{0,6}^{\text{ord}} &= \frac{1}{3}p(p^2 + 2) \\ n_{1,3}^{\text{ord}} &= \frac{1}{2}p(p^2 + 1) \\ n_{2,0}^{\text{ord}} &= \frac{1}{3}p(2p^2 + 1) \end{aligned}$$

Remark. Just as we defined “discrete models of $\overline{\mathcal{N}}_{g,r}$ at infinity” $T_{0,r}$; $T_{1,r}$ above for the cases $g = 0, 1$, one can, more generally, define such discrete models for arbitrary (g, r) that model the combinatorics involved in Theorem 1.1. Also, by letting p go to infinity, one can obtain from these models “continuous limits” $\mathcal{T}_{0,r}$; $\mathcal{T}_{1,r}$, etc. that estimate the coefficient of the term of highest order (relative to p) of the corresponding discrete models. *It would be interesting to investigate whether these continuous limits have anything to do with (classical) Teichmüller theory over the complex numbers.*

Remark. One interesting aspect of Theorem 1.1 is that, for instance in the case of $n_{g,r}^{\text{ord}}$, Theorem 1.1 gives, for different totally degenerate

r -pointed stable curves of genus g , essentially different ways of computing the same number $n_{g,r}^{\text{ord}}$. If one writes out what this means, one sees that *one obtains various nontrivial combinatorial identities (involving sums like those appearing in Corollary 1.3, (1), (2))*. It is not clear to me what significance these identities might have, but they at least appear to be (in a computational sense) nontrivial.

§2. The Main Theorem

In this §, we develop (in the first three subsections) the technical material that will be used to prove the main result of this Chapter, and then we give a proof (in the last subsection) of the main theorem of this Chapter.

§2.1. The Aphillal Case

In this Chapter, p will always be an odd prime. Let S^{\log} be a fine noetherian log scheme of characteristic p . Let $X^{\log} \rightarrow S^{\log}$ be an r -pointed stable log-curve of genus g . Let (P, ∇_P) be a *nilpotent torally crys-stable bundle* on X^{\log} (of radii ρ). Let $\mathcal{P}^\vee : \text{Ad}(P) \rightarrow \Phi_X^* \omega_{X^{\log}/S^{\log}}$ be the dual of the p -curvature of (P, ∇_P) . Let us write

$$V_{\mathcal{P}} \subseteq X$$

for the closed subscheme which is the zero locus of \mathcal{P}^\vee . Now we would like to make some new definitions:

Definition 2.1. We shall call (P, ∇_P) *conservatively spiked* if the pull-back of $V_{\mathcal{P}}$ to the normalization Z of any irreducible component of a geometric fiber of $X \rightarrow S$ is either all of Z , or a divisor $D \subseteq Z$ such that the order of D at any closed point of Z which maps to a node of X is $< p$.

In this subsection, we will mainly only be concerned with conservatively spiked (P, ∇_P) .

Let $M_f \subseteq X$ be the divisor of marked points. If (P, ∇_P) is, in addition, *torally indigenious*, then one can compose the inclusion

$$f_* \omega_{X^{\log}/S^{\log}}^{\otimes 2}(-M_f) \subseteq \mathbf{R}^1 f_{\text{DR},*} \text{Ad}(P)$$

(induced by the Hodge section of $P \rightarrow X$) with the morphism

$$\Psi_P : \mathbf{R}^1 f_{\mathrm{DR},*} \mathrm{Ad}(P) \rightarrow \Phi_S^*(f_* \omega_{X^{\log}/S^{\log}}^{\otimes 2}(-M_f))$$

induced by \mathcal{P}^\vee and the Cartier operator.

Definition 2.2. Suppose that (P, ∇_P) is torally indigenous. Then if the composite morphism

$$\Phi_P : f_* \omega_{X^{\log}/S^{\log}}^{\otimes 2}(-M_f) \rightarrow \Phi_S^*(f_*(\omega_{X^{\log}/S^{\log}}^{\otimes 2}(-M_f)))$$

just defined has constant rank κ , we shall say that (P, ∇_P) has *toral-ordinary rank* κ . If $\kappa = 3g - 3 + r$, then we shall call (P, ∇_P) *torally ordinary*.

Thus, if (P, ∇_P) is indigenous and torally ordinary, then it is ordinary. Also, note that if (P, ∇_P) is a torally ordinary torally indigenous bundle, then the natural morphism $\overline{\mathcal{N}}_{g,r}^\rho \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ is *étale* at the S -valued point of $\overline{\mathcal{N}}_{g,r}^\rho$ defined by (P, ∇_P) . (Indeed, this follows immediately from the fact that the kernel of the morphism

$$\Psi_P : \mathbf{R}^1 f_{\mathrm{DR},*} \mathrm{Ad}(P) \rightarrow \Phi_S^*(f_* \omega_{X^{\log}/S^{\log}}^{\otimes 2}(-M_f))$$

may be identified with the logarithmic tangent bundle to $(\overline{\mathcal{N}}_{g,r}^\rho)^{\log}$; thus, the condition of “toral ordinariness” implies that the inclusion $\overline{\mathcal{N}}_{g,r}^\rho \subseteq (\overline{\mathcal{S}}_{g,r})_{\mathbf{F}_p}$ is transversal to the projection $(\overline{\mathcal{S}}_{g,r})_{\mathbf{F}_p} \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ at the S -valued point defined by (P, ∇_P) .)

Now suppose that (P, ∇_P) is a *conservatively spiked* nilpotent torally crys-stable bundle (i.e., not necessarily torally indigenous). Let k be an algebraically closed field. Let $\nu \in X(k)$ be a node lying over a point $s \in S(k)$. Let Z_1 and Z_2 be the irreducible components of the formal completion of the fiber X_s at ν . Let $D_i \subseteq Z_i$ be the restriction of $V_{\mathcal{P}}$ to Z_i , for $i = 1, 2$.

Definition 2.3. We shall say that (P, ∇_P) is *aphilial* at ν if D_i is a divisor of order d_i on Z_i (for $i = 1, 2$) with $d_1 = d_2 \neq 0$. (Note that in particular, we are assuming here that $D_i \neq Z_i$, for $i = 1, 2$.)

In this subsection, we shall mainly be concerned with (P, ∇_P) that are aphilial at certain nodes. Note that we always (even without the assumption of aphiliality) have $d_1 \equiv \pm d_2 \pmod{p}$ (cf. the discussion of Chapter II, §3.1). Thus, the aphilial case is one of the main cases *not* dealt with in the deformation theory of Chapter II, §2 and 3. Indeed, the “mildly spiked case” of *loc. cit.* is precisely the case where $d_1 + d_2 = p$ (cf. Chapter II, §3.1).

Remark. The reason for the term “aphilial” (Greek for “unfriendly”) is that unlike the “mildly spiked case,” where D_1 and D_2 fuse under deformation to form a single irreducible divisor of degree p (as was seen in Chapter II, §3), in the aphilial case, we shall see that deformation annihilates D_1 and D_2 , i.e., there are no points of $V_{\mathcal{P}}$ that specialize to D_1 or D_2 .

Now let us fix a node $\nu \in S(k)$ at which (P, ∇_P) is *aphilial*. Since we wish to study how aphilial nodes deform, we will be interested in the case where the base S is a trait. Thus, we assume that $S = \text{Spec}(A)$, where A is a complete discrete valuation ring whose residue field is k . Moreover, we assume that the log structure on S is that defined by the special point, and that the generic fiber of $X \rightarrow S$ is *smooth*. We would like to study $V_{\mathcal{P}}$ in a neighborhood of ν . First, let us note the following

Lemma 2.4. *After possibly replacing S by a finite extension of S , there exists a pointed semi-stable log-curve $Y^{\log} \rightarrow S^{\log}$, together with an S^{\log} -morphism $\phi^{\log} : Y^{\log} \rightarrow X^{\log}$ such that*

- (1) ϕ^{\log} is an isomorphism over the complement of ν in X .
- (2) Y is regular at all the points of $\phi^{-1}(\nu)$.
- (3) $\phi^{-1}(\nu)$ consists of an odd number of irreducible components.
- (4) Write $V_Y \subseteq Y$ for the pull-back of the closed subscheme $V_{\mathcal{P}} \subseteq X$ to Y . Then V_Y is a vertical divisor at all critical the nodes of Y (i.e., nodes of Y that lie over ν). Here, by “vertical” we mean a divisor which set-theoretically lies in the special fiber Y_s of Y over S .

Remark. Thus, in particular, *a priori*, V_Y may have subschemes which are S -flat (i.e., “horizontal”) divisors, but such divisors are assumed not to intersect any critical nodes of Y . Also, away from the critical nodes of Y , V_Y may have connected components or embedded primes of codimension 2. The principal aim of the rest of this subsection is to show, however, that in fact, none of these phenomena actually occur at points of $\phi^{-1}(\nu)$.

Proof. Note that once S is fixed, conditions (1) and (2) uniquely determine $\phi^{\log} : Y^{\log} \rightarrow X^{\log}$. Condition (3) may be satisfied by enlarging S . Thus, the only nontrivial condition is condition (4). Fix

some $\phi_0^{\log} : Y_0^{\log} \rightarrow X^{\log}$ such that conditions (1), (2), and (3) are satisfied. Note that $\phi_0^{-1}(\nu)$ is a chain C_0 of \mathbf{P}^1 's. If we blow-up Y_0 at a point which is a "node" in this chain C_0 of \mathbf{P}^1 's, then we get some Y_1 such that the set-theoretic inverse image of C_0 is some chain $C_1 \subseteq Y_1$. By repeatedly blowing up at nodes of the chain of \mathbf{P}^1 's which is the set-theoretic inverse image of $C_0 \subseteq Y_0$, we get a sequence of blow-ups $\dots (Y_i, C_i) \rightarrow \dots \rightarrow (Y_1, C_1) \rightarrow (Y_0, C_0)$. Moreover, it follows from the theory of resolution of indeterminacy loci on surfaces (see, e.g., Theorem 5.5 of Chapter V of [Harts2]) that we may arrange so that for some i , the (schematic) inverse image of $V_{\mathcal{P}}$ in Y_i is a vertical divisor at each of the nodes of C_i .

Next, let us observe that if R is the completion of Y_i at some point of C_i , then R is a regular local ring of dimension 2 with a regular system of parameters x, y such that $x^a \cdot y^b = t$, for some uniformizer t of A , and nonnegative integers a, b . (Indeed, this follows from a simple calculation of blow-ups.) Now let $A' \stackrel{\text{def}}{=} A[t^{\frac{1}{\gamma}}]$; $S' \stackrel{\text{def}}{=} \text{Spec}(A')$, where γ is a positive integer that divides all the integers a, b that occur in this way. Thus, S' is a strictly henselian trait, so it follows that any *nonnodal* k -valued point of C_i lifts to an A' -valued point of Y_i . Let $Y' \rightarrow S'$ be the unique regular semi-stable curve which is equal to $X \times_{S'} S'$ away from ν . Then A' -valued points of Y_i define A' -valued points of Y' which avoid the nodes of Y' over ν (since Y' is *regular*). It thus follows that:


The natural rational map from Y' to Y_i extends to a morphism from Y' to Y_i at the nodes of Y' over ν which maps the nodes of Y' to nodes of $C_i \subseteq Y_i$.

Now recall that we know that the pull-back of $V_{\mathcal{P}}$ to Y_i is a vertical divisor at the nodes of $C_i \subseteq Y_i$. Thus, it follows that the pull-back of $V_{\mathcal{P}}$ to Y' is a vertical divisor at the critical nodes of Y' , as desired. \circ

Let

$$\phi^{\log} : Y^{\log} \rightarrow X^{\log}$$

be as in Lemma 2.4. Let us denote by Z_1, \dots, Z_N (where N is odd) the (distinct) irreducible components, in order, of $\phi^{-1}(\nu)$. Thus, all the Z_i are isomorphic to \mathbf{P}_k^1 , and they are arranged in a chain, which proceeds (from left to right) Z_1, \dots, Z_N . For $i = 1, \dots, N-1$, let ν_i be the node that joins Z_i to Z_{i+1} . Let ν_0 (respectively, ν_N) be the node on Z_1 (respectively, Z_N) that meets an irreducible component of Y_s which is not among the Z_i . Let Z_0 (respectively, Z_{N+1}) be the completion at ν_0 (respectively, ν_N) of the irreducible component of Y_s which is not among the Z_i . Thus, $\phi^{-1}(\nu)$ may be envisioned as follows:



Now let

$$Z_\infty \stackrel{\text{def}}{=} \sum_{i=1}^N c_i Z_i \subseteq V_Y$$

be the *maximal* (with respect to the obvious ordering by inclusion) divisor contained scheme-theoretically in V_Y , and set-theoretically in $\phi^{-1}(\nu)$. Thus, all the c_i 's are *positive integers*. Let \mathcal{L} be the restriction of the line bundle $(\Phi_Y^* \omega_{Y/\log/S^{\log}})(-Z_\infty)$ to Y_s . Let \mathcal{E} be the restriction of $\text{Ad}(P)$ to Y_s . Then \mathcal{P}^\vee (the dual of the p -curvature of (P, ∇_P)) defines a morphism of vector bundles

$$\xi: \mathcal{E} \rightarrow \mathcal{L}$$

which is generically nonzero (by maximality) on each of the Z_i (for $i = 0, \dots, N+1$). Moreover, by (4) of Lemma 2.4, ξ is surjective at all the critical nodes of Y_s .

Next observe that \mathcal{E} restricts to the trivial vector bundle of rank 3 on each Z_i (for $i = 1, \dots, N$). It thus follows that for $i = 1, \dots, N$, $\deg(\mathcal{L}|_{Z_i}) \geq 0$. On the other hand, one computes easily that (for $i = 1, \dots, N$) $\deg(\mathcal{L}|_{Z_i}) = 2c_i - c_{i-1} - c_{i+1}$. Thus, we obtain the following

Observation. For $i = 1, \dots, N$, we have: $2c_i \geq c_{i-1} + c_{i+1}$. (Here, for convenience, we set $c_0 = c_{N+1} = 0$.)

Let a be the order of the divisor $V_Y|_{Z_0}$. By the definition of aphillity, it follows that a is also the order of the divisor $V_Y|_{Z_{N+1}}$. By (4) of Lemma 2.4, it follows, moreover, that $c_1 = c_N = a$. Now we are ready to prove the following result:

Theorem 2.5. If $1 \leq i \leq \frac{1}{2}(N+1)$, then $c_i = i \cdot a$. If $\frac{1}{2}(N+1) \leq i \leq N$, then $c_i = (N+1-i)a$. Moreover, V_Y and Z_∞ coincide scheme-theoretically in a neighborhood of $\phi^{-1}(\nu)$.

Proof. We shall prove the first statement by induction on i . The second statement follows immediately from the first by left-right symmetry. Thus, let $2 \leq i \leq \frac{1}{2}(N+1)$, and suppose that we know that $c_j = j \cdot a$, for all $j \in \{1, \dots, i-1\}$. By left-right symmetry, we may also assume that we know that $c_j = (N+1-j)a$, for all $j \in \{N+1-(i-1), \dots, N\}$. Thus, in particular, we know that $c_{i-1} = c_{N+2-i} = (i-1)a$. Then I claim that $c_i \geq (i-1)a$. Indeed, to verify this claim, one simply considers the differences $\delta_j \stackrel{\text{def}}{=} c_j - c_{j-1}$ (for $j = i, \dots, N+2-i$): By the Observation preceding Theorem 2.5, one knows that δ_j is *monotone decreasing* as a function of j . Thus, if δ_i were negative, all the δ_j in this range would have to be negative, so (adding things up) we would obtain that $(i-1)a = c_{N+2-i} < c_{i-1} = (i-1)a$, which is absurd. Thus, $c_i - (i-1)a = c_i - c_{i-1} = \delta_i \geq 0$, as claimed. On the other hand, by the Observation preceding Theorem 2.5, we have that $c_i + (i-2)a = c_i + c_{i-2} \leq 2c_{i-1} = (2i-2)a$. Thus, $c_i \leq i \cdot a$. In other words, so far we have proven that $(i-1)a \leq c_i \leq i \cdot a$.

Let $Z \stackrel{\text{def}}{=} Z_{i-1}$. The next step is to study the restrictions \mathcal{E}_Z and \mathcal{L}_Z of \mathcal{E} and \mathcal{L} to Z . Note first of all that the given connections on $\text{Ad}(P)$ and $\Phi_Y^* \omega_{Y^{\log}/S^{\log}}$ induce connections $\nabla_{\mathcal{E}}$ and $\nabla_{\mathcal{L}}$ on \mathcal{E} and \mathcal{L} relative to which ξ is *horizontal*. Moreover, since $\phi(Z) = \{\nu\}$, and \mathcal{E} is obtained by pull-back via $\phi: Y \rightarrow X$, it follows that the decomposition of $\text{Ad}(P)|_{\nu}$ into eigenspaces for the monodromy operator at ν induces the following decomposition (compatible with connections) on $(\mathcal{E}_Z, \nabla_{\mathcal{E}|_Z})$:

$$\mathcal{E}_Z \cong \mathcal{A} \oplus \mathcal{O}_Z \oplus \mathcal{B}$$

Here \mathcal{O}_Z is equipped with the trivial connection; $\mathcal{A} = \mathcal{O}_Z(a' \cdot \nu_{i-2} - a' \cdot \nu_{i-1})$ (where a' is a nonnegative integer $< p$, and the symbols “ ν_{i-2} ” and “ ν_{i-1} ” stand for the divisors on Z defined by the points ν_{i-2} and ν_{i-1} , respectively) equipped with the logarithmic connection induced by the trivial connection on \mathcal{O}_Z ; and $(\mathcal{B}, \nabla_{\mathcal{B}})$ is the dual of $(\mathcal{A}, \nabla_{\mathcal{A}})$. Note, moreover, that \mathcal{E}_Z has a Lie algebra structure induced by the Lie algebra structure of $\text{Ad}(P)$. Relative to this Lie algebra structure, the direct summands \mathcal{A} and \mathcal{B} are nilpotent; \mathcal{O}_Z is semi-simple; and \mathcal{O}_Z stabilizes \mathcal{A} and \mathcal{B} (under the bracket operation).

Next let us consider \mathcal{L}_Z . Write $(\mathcal{M}, \nabla_{\mathcal{M}})$ for $(\mathcal{L}_Z, \nabla_{\mathcal{L}|_Z})$. Then note that we can identify \mathcal{M} with $\mathcal{O}_Z(n \cdot \nu_{i-2} + m \cdot \nu_{i-1})$, where $n = c_{i-1} - c_{i-2} = a$; and $m = c_{i-1} - c_i = (i-1)a - c_i$. Moreover, the connection $\nabla_{\mathcal{M}}$ is precisely the logarithmic connection induced on $\mathcal{O}_Z(n \cdot \nu_{i-2} + m \cdot \nu_{i-1})$ by the trivial connection on \mathcal{O}_Z . Thus, $(\mathcal{M}, \nabla_{\mathcal{M}})$ has monodromy equal to $-a$ (respectively, $-m$) at ν_{i-2} (respectively, ν_{i-1}). Finally, recall that we are interested in determining m , i.e., showing that $m = -a$. So far, we know that $0 \geq m \geq -a$.

To obtain more information, we must consider the morphism $\xi_Z: \mathcal{E}_Z \rightarrow \mathcal{M}_Z$ in greater detail. First, observe that since ξ_Z is surjective

at ν_{i-2} and ν_{i-1} and the dual to ξ_Z has nilpotent image, it follows (by considering the action of the monodromy operators at ν_{i-2} and ν_{i-1}) that

$$m \equiv \pm a' \equiv \pm a \pmod{p}$$

It thus follows that m is equal to either $a - p$ or $-a$. Also, we may assume (by switching \mathcal{A} and \mathcal{B}) that $a' = a$.

Now consider the morphism $\xi_{\mathcal{O}} : \mathcal{O}_Z \subseteq \mathcal{E}_Z \rightarrow \mathcal{M}$ (induced by restricting ξ_Z to the middle direct summand of \mathcal{E}_Z). By monodromy considerations, it thus follows that $\xi_{\mathcal{O}}$ in fact maps into $\mathcal{O}_Z(m \cdot \nu_{i-1})$ (a line bundle of *negative* degree since our two candidates for m , namely, $a - p$ and $-a$, are both < 0). Thus, we see that $\xi_{\mathcal{O}}$ must be identically zero. On the other hand, by nilpotence, it then follows that ξ_Z must factor through either the projection to \mathcal{A} or the projection to \mathcal{B} . By monodromy considerations at ν_{i-2} (and the fact that $a \not\equiv -a \pmod{p}$), it thus follows that ξ_Z could never be surjective at ν_{i-2} unless ξ_Z factors through the projection to \mathcal{A} . We thus conclude that ξ_Z factors through the projection to \mathcal{A} , and that $(\mathcal{A}, \nabla_{\mathcal{A}}) \cong (\mathcal{M}, \nabla_{\mathcal{M}})$. By considering the monodromy at ν_{i-1} , we thus obtain that $m = -a$, so $c_i = i \cdot a$, as desired. This completes the proof of the first two statements of the Theorem.

Finally, we turn to the proof of the third statement of the Theorem. Note that this third statement is equivalent to showing that $\xi : \mathcal{E} \rightarrow \mathcal{L}$ is surjective in a neighborhood of $\phi^{-1}(\nu)$. Moreover, it follows from what we have already done that ξ is surjective over all the Z_i except possibly $Z_{\frac{1}{2}(N+1)}$. Thus, let us now set $Z \stackrel{\text{def}}{=} Z_{\frac{1}{2}(N+1)}$. Then just as before, we have a decomposition $\mathcal{E}_Z \cong \mathcal{A} \oplus \mathcal{O}_Z \oplus \mathcal{B}$; and $\mathcal{M} \stackrel{\text{def}}{=} \mathcal{L}_Z = \mathcal{O}_Z(a \cdot \nu_{\frac{1}{2}(N-1)} + a \cdot \nu_{\frac{1}{2}(N+1)})$. Now let us consider the restriction $\xi_{\mathcal{O}} : \mathcal{O}_Z \subseteq \mathcal{E}_Z \rightarrow \mathcal{M}$ of ξ_Z to the middle factor \mathcal{O}_Z of \mathcal{E}_Z . By monodromy considerations, it is clear that $\xi_{\mathcal{O}}$ must factor through $\mathcal{O}_Z \subseteq \mathcal{M}$. Thus, *either* $\xi_{\mathcal{O}}$ is identically zero, *or* it is nonzero everywhere away from the nodes. But if it were identically zero, then we would obtain a contradiction as in the preceding paragraph (i.e., ξ_Z would have to factor through the projection to \mathcal{A} or the projection to \mathcal{B} , but since $a \not\equiv -a \pmod{p}$, this would violate the surjectivity of ξ_Z at the nodes). Thus, we conclude that $\xi_{\mathcal{O}}$ is surjective away from the nodes. This completes the proof of the third statement of the Theorem. \circ

Now, for the rest of this subsection, let us assume that (P, ∇_P) is *totally indigenous*. Let $\psi^{\log} : Z^{\log} \rightarrow X_s^{\log}$ be the partial normalization of X_s^{\log} at the node ν . Here, the log structure on Z^{\log} is the pull-back of that on X_s^{\log} away from $\psi^{-1}(\nu)$, and (in a sufficiently small neighborhood of $\psi^{-1}(\nu)$) equal to that given by regarding the points of $\psi^{-1}(\nu)$ as marked points of Z . Thus, the restriction of (P, ∇_P) to Z^{\log}

is torally indigenous on Z^{\log} . Let us write X_η^{\log} for the generic fiber of $X^{\log} \rightarrow S^{\log}$.

Corollary 2.6. *Suppose that (P, ∇_P) is torally indigenous, and that $(P, \nabla_P)|_{Z^{\log}}$ has toral-ordinary rank κ . Then, $(P, \nabla_P)|_{X_\eta^{\log}}$ is toral-ordinary rank $\geq \kappa + 1$.*

Proof. Let us denote the irreducible component $Z_{\frac{1}{2}(N+1)}$ of $\phi^{-1}(\nu)$ by Z . We claim that $\Phi_P(\alpha)|_Y$ (regarded as a section of the line bundle $\Phi_S^*(\omega_{Y^{\log}/S^{\log}})^{\otimes 2}$) has a zero of order precisely $\frac{1}{2}(N+1)a$ at the divisor Z of Y . Before proving this claim, let us first observe that it implies that Φ_P generically has rank $\geq \kappa + 1$. Indeed, let us consider the morphism

$$\Phi'_P : H^0(X, \omega_{X^{\log}/S^{\log}}^{\otimes 2}(-M_f)) \rightarrow \Phi_S^* H^0(Y, \omega_{Y^{\log}/S^{\log}}^{\otimes 2}(-M_f - Z_\infty))$$

(which is defined similarly to Φ_P , and may, in fact, be identified with Φ_P over the generic point of S). To show that Φ_P generically has rank $\geq \kappa + 1$, it thus suffices to show that Φ'_P generically has rank $\geq \kappa + 1$. But to show that Φ'_P generically has rank $\geq \kappa + 1$, it suffices to show that $\Phi'_P \otimes k$ has rank $\geq \kappa + 1$. To do this, we observe that the rank of the image \mathcal{I} under $\Phi'_P \otimes k$ of the sections of $H^0(X_s, \omega_{X_s^{\log}/S^{\log}}^{\otimes 2} \otimes k)$ that vanish at ν is $\geq \kappa$ (since $(P, \nabla_P)|_{Z^{\log}}$ has toral-ordinary rank κ). Moreover, this image \mathcal{I} clearly vanishes on Z (since $Z \subseteq Y$ maps to ν , and the elements of \mathcal{I} arise from sections that vanish at ν). On the other hand, the validity of the claim implies that $\Phi'_P(\alpha) \otimes k$ (for α as in the claim) is nonzero on Z . Thus, $\Phi'_P \otimes k$ has rank $\geq \kappa + 1$, as desired.

It remains to prove the claim. This amounts to computing in greater detail the morphism ξ_Z considered in the last paragraph of the proof of Theorem 2.5. We use the notation of *loc. cit.* Also, let us write $[0]$ (respectively, $[\infty]$) for $\nu_{\frac{1}{2}(N-1)}$ (respectively, $\nu_{\frac{1}{2}(N+1)}$). Now observe that monodromy and degree considerations imply that the morphism

$$\xi_A : \mathcal{A} = \mathcal{O}_Z(a \cdot [0] - a \cdot [\infty]) \rightarrow \mathcal{M} = \mathcal{O}_Z(a \cdot [0] + a \cdot [\infty])$$

(where ξ_A is ξ_Z restricted to the factor $\mathcal{A} \subseteq \mathcal{E}_Z$) is a k -linear combination of multiplication by 1 and t^p (where t is the standard rational function on $Z \cong \mathbf{P}_k^1$ which is 0 at $[0]$ and has a pole at $[\infty]$). Thus, the image of the section t^{-a} of \mathcal{A} is a k -linear combination of t^{-a} and t^{p-a} . Moreover, if we apply the Cartier operator to t^{-a} or t^{p-a} (multiplied by $\frac{dt}{t}$, which is the section of \mathcal{M} used to identify \mathcal{M} with $\mathcal{O}_Z(a \cdot [0] + a \cdot [\infty])$), we get 0. That is to say, we have proven that:

If we apply ξ_Z , followed by the Cartier operator, to any horizontal section of $\Gamma(Z, \mathcal{A}) \subseteq \Gamma(Z, \mathcal{E}_Z)$, we get zero.

A similar result holds for ξ_B .

On the other hand, the last paragraph of the proof of Theorem 2.5 implies that the image of the section 1 of \mathcal{O}_Z under the morphism

$$\xi_{\mathcal{O}} : \mathcal{O}_Z \rightarrow \mathcal{M} = \mathcal{O}_Z(a \cdot [0] + a \cdot [\infty])$$

(where $\xi_{\mathcal{O}}$ is ξ_Z restricted to the factor $\mathcal{O}_Z \subseteq \mathcal{E}_Z$) is a nonzero k -multiple of the section 1 of \mathcal{M} . Moreover, if we apply the Cartier operator to $\frac{dt}{t}$, we get $\frac{dt}{t}$.

Now we are ready to compute $\Phi'_P(\alpha) \otimes k$. Since α is nonzero at ν , the restriction of α to $\text{Ad}(P)|_Z = \mathcal{E}_Z$ is nonzero. Moreover, since the Hodge filtration is not horizontal (i.e., since (P, ∇_P) is torally *indigenous*), it follows that the “ \mathcal{O}_Z -component” of the restriction of α to \mathcal{E}_Z is *nonzero*. The computations of the preceding two paragraphs thus imply that $\Phi'_P(\alpha) \otimes k$ depends only on this \mathcal{O}_Z -component, and, moreover, that since this \mathcal{O}_Z -component is nonzero, $\Phi'_P(\alpha) \otimes k$ itself is nonzero at the generic point of Z , as desired. This completes the proof of the claim and hence of Corollary 2.6. \circ


§2.2. Grafting on Dormant Atoms I: Virtual p-Curvatures

The purpose of this and the following subsection is to study the deformations of a stable nilcurve obtained by concatenating – “grafting” – a dormant atom to a mildly spiked (Chapter II, §3) nilcurve. Just as in the case of mildly spiked nilcurves, this will consist first of performing various calculations on the initial given stable nilcurve over a field (which we do in this subsection), and then applying these calculations to the deformation theory (which we do in the following subsection).

Thus, let S^{\log} be a fine log scheme whose underlying scheme is $\text{Spec}(k)$, where k is an algebraically closed field of characteristic p . Let $X^{\log} \rightarrow S^{\log}$ be an r -pointed stable log-curve of genus g . Let us assume that X^{\log} is obtained by gluing together an $(r-1)$ -pointed stable log-curve of genus g , $Y^{\log} \rightarrow S^{\log}$, to a 3-pointed (smooth) log-curve of genus 0, $Z^{\log} \rightarrow S^{\log}$, at a single marked point, which we shall henceforth refer to as the *critical point* of X^{\log} (or Y^{\log} or Z^{\log}). Let us refer to the marked points of Z via the symbols α , β , and γ , where α is the critical point. Below, we provide an illustration of the curve X^{\log} .

Let us assume next that we are given a *dormant* torally indigenous bundle (R, ∇_R) on Z^{\log} , of diameters a , b , and c (which we shall think

of as positive integers $\leq p-1$) at the points α , β , and γ , respectively. Without loss of generality, we may assume that $a+b+c$ is odd.



Let us regard Z as $\mathbf{P}(V)$, where V is some two-dimensional k -vector space. Write $V_\alpha \subseteq V$ (respectively, $V_\beta \subseteq V$; $V_\gamma \subseteq V$) for the one-dimensional k -subspaces corresponding to the respective marked points. Let \mathcal{F} be the rank two vector bundle on Z given by taking $V \otimes_k \mathcal{O}_Z$ and modifying the integral structure at α by taking the direct limit of the following diagram:

$$\begin{array}{ccc} V_\alpha \otimes_k \mathcal{O}_Z & \longrightarrow & V \otimes_k \mathcal{O}_Z \\ \downarrow t_\alpha^* & & \\ V_\alpha \otimes_k \mathcal{O}_Z & & \end{array}$$

(where t_α is a local coordinate on Z that vanishes at α ; the horizontal morphism is the natural inclusion); moreover, we modify the integral structure in a similar fashion at β and γ . Then one knows from the theory of dormant atoms (Chapter IV, §2.1, Theorem 2.3) that R may be identified with $\mathbf{P}(\mathcal{F})$. (Note, however, that the Hodge section of R is *not* induced by the usual Hodge section of $\mathbf{P}(V)$. Rather, it is obtained by pulling back by some morphism “ ψ ” (cf. Chapter IV, Theorem 2.3) the usual Hodge section of the trivial \mathbf{P}^1 -bundle on another (i.e., $\neq Z$) copy of \mathbf{P}^1 .)

Let $\sigma_\beta : Z \rightarrow R = \mathbf{P}(\mathcal{F})$ be the section defined by $V_\beta \otimes_k \mathcal{O}_Z \subseteq V \otimes_k \mathcal{O}_Z$. Similarly, we have $\sigma_\gamma : Z \rightarrow R$. Note that both σ_β and σ_γ are *horizontal sections* of R (relative to ∇_R), and, moreover, the canonical height of σ_β (respectively, σ_γ) is equal to $\frac{1}{2}(a+c-b)$ (respectively, $\frac{1}{2}(a+b-c)$). Finally, observe that σ_β (respectively, σ_γ) defines a horizontal surjection

$$\mathrm{Ad}(R) \rightarrow \mathcal{V}_\beta \quad (\text{respectively, } \mathrm{Ad}(R) \rightarrow \mathcal{V}_\gamma)$$

whose dual has nilpotent image in $\mathrm{Ad}(R)$. Thus, \mathcal{V}_β (respectively, \mathcal{V}_γ) inherits a connection $\nabla_{\mathcal{V}_\beta}$ (respectively, $\nabla_{\mathcal{V}_\gamma}$), necessarily of p -curvature zero, from ∇_R .

Definition 2.7. We shall refer to the quotients $\mathrm{Ad}(R) \rightarrow \mathcal{V}_\beta$ and $\mathrm{Ad}(R) \rightarrow \mathcal{V}_\gamma$ as *virtual p -curvatures on Z^{\log}* .

The reason for the name “virtual p -curvature” will become obvious after we study the deformation theory in the following subsection. Roughly speaking, the “virtual p -curvature” arises as the result of specializing the p -curvature of some deformation of a torally indogenous bundle on X^{\log} whose restriction to Z^{\log} is (R, ∇_R) . Note that the two virtual p -curvatures are always *distinct*.

For the rest of the discussion in this subsection, we would like to fix a single virtual p -curvature $\mathrm{Ad}(R) \rightarrow \mathcal{V}$. Without loss of generality, (in order to fix ideas) let us assume that we have chosen the virtual p -curvature $\mathrm{Ad}(R) \rightarrow \mathcal{V}_\beta$. Thus, \mathcal{V} has degree $a+c-b$, and the monodromy of $(\mathcal{V}, \nabla_\mathcal{V})$ at α (respectively, β ; γ) is equal to $-a$ (respectively, b ; $-c$). In other words, we can identify $(\mathcal{V}, \nabla_\mathcal{V})$ with the line bundle $\mathcal{O}_Z(a \cdot \alpha - b \cdot \beta + c \cdot \gamma)$, equipped with the logarithmic connection induced by the trivial connection on \mathcal{O}_Z . In particular, it follows that \mathcal{V} has *no global horizontal sections over Z* .

Now let us suppose that we are given a *nilpotent torally crys-stable bundle* (P, ∇_P) on X^{\log} whose restriction to Z^{\log} is (R, ∇_R) . Let us denote its restriction to Y^{\log} by (Q, ∇_Q) . That is, we have:

$$(P, \nabla_P)|_{Y^{\log}} = (Q, \nabla_Q); \quad (P, \nabla_P)|_{Z^{\log}} = (R, \nabla_R)$$

Let us suppose, moreover, that (Q, ∇_Q) is *mildly spiked* (Chapter II, §3.1), and that

The p -curvature of (Q, ∇_Q) vanishes to order precisely a at the critical point of Y .

Let $\mathrm{Ad}(Q) \rightarrow \mathcal{U}$ be the quotient induced by taking the image of the dual to the p -curvature of (Q, ∇_Q) . Then it follows from the assumption that (Q, ∇_Q) is mildly spiked that \mathcal{U} is a line bundle on Y . Moreover, the monodromy of \mathcal{U} at the point α is a , i.e., precisely minus the monodromy of $(\mathcal{V}, \nabla_\mathcal{V})$ at α . Thus, it follows that the quotients $\mathrm{Ad}(Q) \rightarrow \mathcal{U}$ and $\mathrm{Ad}(R) \rightarrow \mathcal{V}$ glue together to form a quotient

$$\mathrm{Ad}(P) \rightarrow \mathcal{W}$$

where \mathcal{W} is also a line bundle. Moreover, \mathcal{W} inherits a connection $\nabla_\mathcal{W}$ from $\mathrm{Ad}(P)$. We would now like to study this quotient $\mathrm{Ad}(P) \rightarrow \mathcal{W}$ from a point of view similar to that of Chapter II, §3.2. (Beware, however, that the notation “ Z ” has a different meaning here from its meaning in *loc. cit.*!)

Let us denote by $\mathrm{Ad}^q(P) \subseteq \mathrm{Ad}(P)$ the subsheaf of sections whose values at the marked points of X^{\log} are orthogonal (relative to the Killing form) to the monodromy operator (cf. Chapter I, §1.4). Then we define the de Rham cohomology of $\mathrm{Ad}(P)$ to be the hypercohomology (on X) of the complex $\mathrm{Ad}(P) \rightarrow \mathrm{Ad}^q(P) \otimes_{\mathcal{O}_X} \omega_{X^{\log}/S^{\log}}$ defined by ∇_P . Let us refer (as in Chapter II, §3.2) to the points of X where $(\mathcal{W}, \nabla_{\mathcal{W}})$ has nonzero monodromy at \mathcal{W} -active. Thus, the \mathcal{W} -active points are either marked points or nodes. Let us denote the divisor of marked points that are \mathcal{W} -active (respectively, not \mathcal{W} -active) by $M_X^{\mathrm{ac}} \subseteq X$ (respectively, $M_X^{\mathrm{non}} \subseteq X$). Let us define the de Rham cohomology of \mathcal{W} to be the hypercohomology of the complex

$$\mathcal{W} \rightarrow \mathcal{W} \otimes_{\mathcal{O}_X} \omega_{X^{\log}/S^{\log}}(-M_X^{\mathrm{non}})$$

defined by $\nabla_{\mathcal{W}}$. Finally, let us denote the image of (the natural morphism)

$$H_{\mathrm{DR}}^1(X^{\log}, \mathrm{Ad}(P)) \rightarrow H_{\mathrm{DR}}^1(X^{\log}, \mathcal{W})$$

by \mathcal{I}_X .

Let $\mathcal{H} \subseteq \mathrm{Ad}(P)$ be the kernel of the surjection $\mathrm{Ad}(P) \rightarrow \mathcal{W}$. Thus, \mathcal{H} fits into an exact sequence

$$0 \rightarrow \mathcal{W}^{\vee} \rightarrow \mathcal{H} \rightarrow \mathcal{O}_X \rightarrow 0$$

Moreover, \mathcal{H} inherits a connection $\nabla_{\mathcal{H}}$ from that of $\mathrm{Ad}(P)$. Let us define the de Rham cohomology of \mathcal{H} to be the hypercohomology of the complex $\mathcal{H} \rightarrow (\mathcal{H} \cap \mathrm{Ad}^q(P)) \otimes_{\mathcal{O}_X} \omega_{X^{\log}/S^{\log}}$.

Lemma 2.8. *The natural projection $H_{\mathrm{DR}}^2(X^{\log}, \mathcal{H}) \rightarrow (H^0(X, \mathcal{O}_X(-M_X^{\mathrm{non}})))^{\vee}$ (induced by $\mathcal{H} \rightarrow \mathcal{O}_X$ and Serre duality) is an isomorphism.*

Proof. By duality, it suffices to verify that the horizontal sections of \mathcal{H}^{\vee} over X actually lie inside $\mathcal{O}_X \subseteq \mathcal{H}^{\vee}$. Over Y , this follows from the fact that the p -curvature is generically nonzero on Y . Over Z , this follows from the fact (noted above) that ν has no global horizontal sections. \circ

Let \mathcal{K}_X (respectively, \mathcal{C}_X) be the kernel (respectively, cokernel) of the morphism $\mathcal{W} \rightarrow \mathcal{W} \otimes_{\mathcal{O}_X} \omega_{X^{\log}/S^{\log}}(-M_X^{\mathrm{non}})$ defined by $\nabla_{\mathcal{W}}$. Then we have a natural composite morphism

$$\zeta_X : H^1(X, \mathcal{K}_X) \hookrightarrow H_{\mathrm{DR}}^1(X^{\log}, \mathcal{W}) \rightarrow H_{\mathrm{DR}}^2(X^{\log}, \mathcal{H})$$

where the first morphism in the composite is the natural one induced by considering the spectral sequence for the complex that defines the de Rham cohomology of \mathcal{W} ; and the second morphism in the composite is the connecting morphism associated to $0 \rightarrow \mathcal{H} \rightarrow \text{Ad}(P) \rightarrow \mathcal{W} \rightarrow 0$.

Now let us take a closer look at \mathcal{K}_X and \mathcal{C}_X . Note that \mathcal{K}_X is (naturally) a direct sum of sheaves \mathcal{K}_Y and \mathcal{K}_Z supported on Y and Z , respectively, i.e.,

$$\mathcal{K}_X = \mathcal{K}_Y \oplus \mathcal{K}_Z$$

Similarly,

$$\mathcal{C}_X = \mathcal{C}_Y \oplus \mathcal{C}_Z$$

Moreover, it follows from the fact that

$$(\mathcal{V}, \nabla_{\mathcal{V}}) = (\mathcal{O}_Z(a \cdot \alpha - b \cdot \beta + c \cdot \gamma), \text{natural } \nabla)$$

that \mathcal{K}_Z is a line bundle on Z^F of degree -1 , while $\mathcal{C}_Z = \mathcal{K}_Z \otimes_{\mathcal{O}_{Z^F}} \omega_{Z^{\log}/S^{\log}}^F$. Thus, \mathcal{C}_Z is a line bundle of degree zero on Z^F . In particular,

$$H^1(Z^F, \mathcal{K}_Z) = 0; \quad \dim_k(H^0(Z^F, \mathcal{C}_Z)) = 1$$

On the other hand, note that by applying the theory of Chapter II, §3.2, *directly to the mildly spiked bundle* (Q, ∇_Q) on Y^{\log} , we obtain a corresponding \mathcal{I}_Y , and $\text{Ker}(\zeta_Y)$ (i.e., the objects denoted simply “ \mathcal{I} ” and “ $\text{Ker}(\zeta)$ ” in *loc. cit.*). Moreover, since $H^1(X, \mathcal{K}_X) = H^1(Y, \mathcal{K}_Y) \oplus H^1(Z, \mathcal{K}_Z) = H^1(Y, \mathcal{K}_Y)$, and

$$H_{\text{DR}}^2(X^{\log}, \mathcal{H}) \cong (H^0(X, \mathcal{O}_X(-M_X^{\text{non}})))^{\vee} = (H^0(Y, \mathcal{O}_Y(-M_Y^{\text{non}})))^{\vee}$$

(note: $M_X^{\text{non}} \subseteq Y \subseteq X$), it follows that ζ_X is “isomorphic” to ζ_Y . In particular, we may identify $\text{Ker}(\zeta_X)$ with $\text{Ker}(\zeta_Y)$.

Proposition 2.9. *There exists a natural exact sequence $0 \rightarrow \text{Ker}(\zeta_X) \rightarrow \mathcal{I}_X \rightarrow H^0(X, \mathcal{C}_X) \rightarrow 0$.*

Proof. The proof is entirely similar to that of Proposition 3.5 of Chapter II: Indeed, the existence of such an exact sequence follows formally from the surjectivity of ζ_X (plus the natural “conjugate exact sequence” $0 \rightarrow H^1(X, \mathcal{K}_X) \rightarrow H_{\text{DR}}^1(X^{\log}, W) \rightarrow H^0(X, \mathcal{C}_X) \rightarrow 0$ for de Rham

cohomology). Moreover, since ζ_X and ζ_Y are “isomorphic,” the surjectivity of ζ_X follows from that of ζ_Y . But the surjectivity of ζ_Y follows from the discussion preceding Proposition 3.5 of Chapter II. \circ

Thus, putting everything together, and applying Proposition 3.5 of Chapter II directly to (Q, ∇_Q) , we obtain the following:

Proposition 2.10. *We have: $\dim_k(\mathcal{I}_X) = \dim_k(\mathcal{I}_Y) + 1$.*

Before proceeding to the deformation theory of the following subsection, we have one more technical detail that we must take care of. First, let \mathcal{H}_Y (respectively, \mathcal{H}_Z) be the restriction of \mathcal{H} to Y (respectively, Z), and let $\mathcal{H}' \subseteq \mathcal{H} \cap \text{Ad}^q(P)$ be the subsheaf consisting of those sections whose image under the projection $\mathcal{H} \rightarrow \mathcal{O}_X$ vanishes at the critical point. Thus, \mathcal{H}' fails to be a vector bundle (only) at the critical point, but instead has the virtue of surjecting naturally onto $\mathcal{H}_Y \cap \text{Ad}^q(Q)$. (Note: Since the critical point becomes a marked point on Y , $\text{Ad}^q(Q) \neq \text{Ad}^q(P)|_Y$.) Moreover, we have an exact sequence

$$0 \rightarrow \mathcal{H}''' \rightarrow \mathcal{H}' \rightarrow \mathcal{H}'' \rightarrow 0$$

where $\mathcal{H}'' \stackrel{\text{def}}{=} \mathcal{H}_Y \cap \text{Ad}^q(Q)$, and $\mathcal{H}''' \stackrel{\text{def}}{=} \mathcal{H}_Z(-\alpha) \cap \text{Ad}^q(R)$.

Next, we would like to consider the commutative diagram:

$$\begin{array}{ccccccc} 0 & \rightarrow & \mathcal{H}_Z(-\alpha) & \rightarrow & \mathcal{H} & \rightarrow & \mathcal{H}_Y & \rightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \rightarrow & \mathcal{H}''' \otimes \omega_{Z^{\log}/S^{\log}} & \rightarrow & \mathcal{H}' \otimes \omega_{X^{\log}/S^{\log}} & \rightarrow & \mathcal{H}'' \otimes \omega_{Y^{\log}/S^{\log}} & \rightarrow & 0 \end{array}$$

Here, the lower exact sequence is obtained by tensoring the exact sequence of the preceding paragraph with $\omega_{X^{\log}/S^{\log}}$. The upper exact sequence is the natural one obtained by restricting sections of \mathcal{H} to Y . The vertical arrows are given by differentiation. We would like to define the de Rham cohomologies of \mathcal{H}''' , \mathcal{H}' , and \mathcal{H}'' to be the hypercohomologies of the first, second, and third vertical arrows of the above diagram. Then it follows from the definitions that we have a long exact cohomology sequence relating the de Rham cohomologies of \mathcal{H}''' , \mathcal{H}' , and \mathcal{H}'' . In particular, we obtain that the cokernel of $H_{\text{DR}}^1(\mathcal{H}') \rightarrow H_{\text{DR}}^1(\mathcal{H}'')$ is isomorphic to the kernel of $H_{\text{DR}}^2(\mathcal{H}''') \rightarrow H_{\text{DR}}^2(\mathcal{H}')$.

Lemma 2.11. *The map $H_{\text{DR}}^2(\mathcal{H}''') \rightarrow H_{\text{DR}}^2(\mathcal{H}')$ is injective.*

Proof. First, let us observe that $H_{\text{DR}}^2(\mathcal{H}''')$ is dual to $H_{\text{DR}}^0(\mathcal{G})$, where

$$\mathcal{G} \stackrel{\text{def}}{=} (\mathcal{H}''')^\vee(-\alpha - \beta - \gamma)$$

Note that \mathcal{H}''' fits into an exact sequence

$$0 \rightarrow \mathcal{V}^\vee(-\alpha) \rightarrow \mathcal{H}''' \rightarrow \mathcal{O}_Z(-\alpha - \beta - \gamma) \rightarrow 0$$

Thus, \mathcal{G} fits into an exact sequence

$$0 \rightarrow \mathcal{O}_Z \rightarrow \mathcal{G} \rightarrow \mathcal{V}(-\beta - \gamma) \rightarrow 0$$

Since, \mathcal{V} has no global horizontal sections, it thus follows that $H_{\text{DR}}^0(\mathcal{G}) = H^0(Z, \mathcal{O}_Z) = k$. Moreover, in order to verify the desired injectivity, it suffices to verify it after further composition with the morphism $H_{\text{DR}}^2(\mathcal{H}') \rightarrow (H_{\text{DR}}^0(Z^{\log}, \mathcal{O}_Z))^\vee$ induced by the surjection $\mathcal{H}' \rightarrow \mathcal{O}_Z(-\alpha - \beta - \gamma)$ and Serre duality. But then it is clear that the resulting map

$$(H_{\text{DR}}^0(Z^{\log}, \mathcal{O}_Z))^\vee = H_{\text{DR}}^2(\mathcal{H}''') \rightarrow H_{\text{DR}}^2(\mathcal{H}') \rightarrow (H_{\text{DR}}^0(Z^{\log}, \mathcal{O}_Z))^\vee$$

is the identity, hence injective. This completes the proof of the Lemma. \square

Corollary 2.12. *The natural morphism*

$$H_{\text{DR}}^1(X^{\log}, \mathcal{H}') \rightarrow H_{\text{DR}}^1(X^{\log}, \mathcal{H}'') = H_{\text{DR}}^1(Y^{\log}, \mathcal{H}_Y)$$

is surjective.

Next, note that we have a natural morphism

$$\begin{aligned} H_{\text{DR}}^1(X^{\log}, \mathcal{H}) &\rightarrow H^0(X, \omega_{X^{\log}/S^{\log}}(-M_X^{\text{ac}}))^F \\ &\rightarrow H^0(Y, \omega_{Y^{\log}/S^{\log}}(-M_Y^{\text{ac}} + \alpha))^F \end{aligned}$$

(where the first morphism is obtained by projecting $\mathcal{H} \rightarrow \mathcal{O}_X$, then applying the map that arises from the “conjugate spectral sequence” that computes the de Rham cohomology of \mathcal{O}_X ; and the second morphism arises from restricting to Y). Similarly, we have a surjection (cf. Chapter II, Proposition 3.6)

$$H_{\text{DR}}^1(Y^{\log}, \mathcal{H}_Y) \rightarrow H^0(Y, \omega_{Y^{\log}/S^{\log}}(-M_Y^{\text{ac}}))^F$$

which we can compose (on the left) with the restriction morphism

$$H_{\mathrm{DR}}^1(X^{\log}, \mathcal{H}') \rightarrow H_{\mathrm{DR}}^1(X^{\log}, \mathcal{H}'') = H_{\mathrm{DR}}^1(Y^{\log}, \mathcal{H}_Y)$$

– which is surjective by Corollary 2.12 – to obtain a surjection

$$H_{\mathrm{DR}}^1(X^{\log}, \mathcal{H}') \rightarrow H^0(Y, \omega_{Y^{\log}/S^{\log}}(-M_Y^{\mathrm{ac}}))^F$$

Thus, we obtain a natural commutative diagram:

$$\begin{array}{ccc} H_{\mathrm{DR}}^1(X^{\log}, \mathcal{H}') & \longrightarrow & H^0(Y, \omega_{Y^{\log}/S^{\log}}(-M_Y^{\mathrm{ac}}))^F \\ \downarrow & & \downarrow \\ H_{\mathrm{DR}}^1(X^{\log}, \mathcal{H}) & \longrightarrow & H^0(Y, \omega_{Y^{\log}/S^{\log}}(-M_Y^{\mathrm{ac}} + \alpha))^F \end{array}$$

Next, we would like to define a subspace $\mathcal{D} \subseteq (H^0(Y, \omega_{Y^{\log}/S^{\log}}(-M_Y^{\mathrm{ac}})))^F$ as follows: Recall that the p -curvature of (Q, ∇_Q) factors through the quotient $\mathrm{Ad}(Q) \rightarrow \mathcal{U}$. Thus, we get a horizontal morphism $\mathcal{U} \rightarrow (\Phi_Y^* \omega_{Y^{\log}/S^{\log}})(-M_Y^{\mathrm{ac}})$. Since this morphism is horizontal, it induces a morphism on the respective subsheaves of horizontal sections, i.e., a morphism

$$\rho : \mathcal{K}_Y \rightarrow (\omega_{Y^{\log}/S^{\log}}(-M_Y^{\mathrm{ac}}))^F$$

hence an injection

$$H^0(\rho) : H^0(Y, \mathcal{K}_Y) \rightarrow (H^0(Y, \omega_{Y^{\log}/S^{\log}}(-M_Y^{\mathrm{ac}})))^F$$

We will denote the image of $H^0(\rho)$ by \mathcal{D} .

Thus, by forming the quotient by \mathcal{D} of the spaces on the right-hand side of the above commutative diagram, we obtain a new commutative diagram

$$\begin{array}{ccc} H_{\mathrm{DR}}^1(X^{\log}, \mathcal{H}') & \longrightarrow & H^0(Y, \omega_{Y^{\log}/S^{\log}}(-M_Y^{\mathrm{ac}}))^F / \mathcal{D} \\ \downarrow & & \downarrow \\ H_{\mathrm{DR}}^1(X^{\log}, \mathcal{H}) & \longrightarrow & H^0(Y, \omega_{Y^{\log}/S^{\log}}(-M_Y^{\mathrm{ac}} + \alpha))^F / \mathcal{D} \end{array}$$

Note that the vertical morphism on the right of this commutative diagram is *injective*, and that (by the preceding discussion) Corollary 2.12 implies that the upper horizontal morphism of this commutative diagram is *surjective*.

In the following subsection, we shall use the above discussion in the form of the following consequence:

Given some $\delta \in H^0(Y, \omega_{Y^{\log}/S^{\log}}(-M_Y^{\text{ac}}))^F/\mathcal{D}$, there always exists an element $\epsilon \in H_{\text{DR}}^1(X^{\log}, \mathcal{H})$ whose image (cf. the preceding commutative diagram) in $H^0(Y, \omega_{Y^{\log}/S^{\log}}(-M_Y^{\text{ac}} + \alpha))^F/\mathcal{D}$ is equal to the image in $H^0(Y, \omega_{Y^{\log}/S^{\log}}(-M_Y^{\text{ac}} + \alpha))^F/\mathcal{D}$ of δ .

§2.3. Grafting on Dormant Atoms II: Deformation Theory

In this subsection, we apply the results of the preceding subsection in a fashion similar to that of Chapter II, §3.3, to conclude certain results concerning the structure of $\bar{N}_{g,r}^\rho$ in a neighborhood of a point corresponding to a nilcurve as in the preceding subsection.

Let $f^{\log} : X^{\log} \rightarrow S^{\log}$ be an r -pointed stable log-curve of genus g , where $2g - 2 + r \geq 1$, and S^{\log} a fine noetherian log scheme of characteristic $p > 2$. Let (P, ∇_P) be a \mathbf{P}^1 -bundle with connection on X^{\log} , of radii ρ at the r marked points of X^{\log} . Let us suppose further that $S = \text{Spec}(A)$, where A is an artinian local ring whose residue field k is algebraically closed. Write $S_0 \subseteq S$ for $\text{Spec}(k) \subseteq \text{Spec}(A)$. Let $I \subseteq A$ be an ideal of length 1. Thus, as an A -module, $I \cong k$. Let us employ a subscript “ I ” (respectively, “0”) to denote the reductions of objects over A modulo I (respectively, the maximal ideal of A). Thus, $S_I = \text{Spec}(A/I)$.

Let us assume that $(P, \nabla_P)_I$ is a *nilpotent torally indigenous bundle* (of radii ρ), and moreover, that $\{X_0^{\log}, (P, \nabla_P)_0\}$ is a nilcurve as in §2.2. Thus, X_0^{\log} is obtained by gluing together some Y_0^{\log} to some Z_0^{\log} (as in §2.2). Moreover, let us assume that we are given a horizontal square nilpotent surjection

$$\kappa_I : \text{Ad}(P_I) \rightarrow \mathcal{W}_I$$

such that $\kappa_0 \stackrel{\text{def}}{=} (\kappa_I)_0$ is the surjection constructed in §2.2 (denoted there by $\text{Ad}(P) \rightarrow \mathcal{W}$) using one of the two virtual p -curvatures. Note that the monodromy of \mathcal{W}_I at the marked points is necessarily the same as that of \mathcal{W}_0 . We shall also assume that κ_I is *compatible with the p -curvature* in the sense that the p -curvature of $(P, \nabla_P)_I$ factors through κ_I . (Note that κ_0 is automatically compatible with the p -curvature.)

We would like to consider the obstruction to lifting $(P, \nabla_P)_I$, together with κ_I , to a nilpotent torally crys-stable bundle of radii ρ over S^{\log} such that the lifting of κ_I remains compatible with the p -curvature. To this end, let us first observe that, by the general nonsense of deformation theory, the obstruction

to lifting κ_I to a horizontal square nilpotent surjection of $\mathrm{Ad}(P)$ (onto some line bundle) forms an element $\eta \in I \otimes_k H_{\mathrm{DR}}^1(X_0^{\log}, \mathcal{W}_0)$ which (tautologically) *vanishes* under the connecting morphism

$$I \otimes_k H_{\mathrm{DR}}^1(X_0^{\log}, \mathcal{W}_0) \rightarrow I \otimes_k H_{\mathrm{DR}}^2(X_0^{\log}, \mathrm{Ker}(\kappa_0))$$

Indeed, this vanishing follows by the argument used in Chapter II, §3.3, plus Lemma 2.8 of the preceding subsection (note that $\mathrm{Ker}(\kappa_0)$ was called “ \mathcal{H} ” in §2.2). Thus, it follows that η lies in $I \otimes_k \mathcal{I}_{X_0}$, i.e., the image of $I \otimes_k H_{\mathrm{DR}}^1(X_0^{\log}, \mathrm{Ad}(P_0))$ in $I \otimes_k H_{\mathrm{DR}}^1(X_0^{\log}, \mathcal{W}_0)$. But this means that by suitably modifying the deformation (P, ∇_P) of $(P, \nabla_P)_I$, we may always assume that there exists a horizontal square nilpotent surjection

$$\kappa' : \mathrm{Ad}(P) \rightarrow \mathcal{W}'$$

whose reduction modulo I is equal to κ_I . Note that the choice of such a κ' is not necessarily unique – indeed, such κ' form a torsor over

$$H_{\mathrm{DR}}^0(X_0^{\log}, \mathcal{W}_0) = H^0(X_0^F, \mathcal{K}_{X_0}) = H^0(Y_0^F, \mathcal{K}_{Y_0})$$

(strictly speaking, tensored over k with I), where $\mathcal{K}_{X_0} \subseteq \mathcal{W}_0$ is the subsheaf of horizontal sections, and \mathcal{K}_{Y_0} is the direct summand of \mathcal{K}_{X_0} which is supported on Y_0 .

Let $\mathcal{B}' = \mathrm{Ad}(P)/\mathrm{Im}((\kappa')^\vee)$. Let $\psi' : \mathrm{Ad}(P) \rightarrow \mathcal{B}'$ be the natural surjection. Note that we have a horizontal exact sequence

$$0 \rightarrow \mathcal{O}_X \rightarrow \mathcal{B}' \rightarrow \mathcal{W}' \rightarrow 0$$

The existence of such a horizontal exact sequence implies that the image of the p -curvature of \mathcal{B}' , hence of (P, ∇_P) , lies in $\mathrm{Ker}(\kappa') \subseteq \mathrm{Ad}(P)$. Thus, the composite of the p -curvature $\mathcal{P} : \Phi_X^* \tau_{X^{\log}/S^{\log}} \rightarrow \mathrm{Ad}(P)$ with κ' must vanish. Unfortunately, however, the composite of the p -curvature $\mathcal{P} : \Phi_X^* \tau_{X^{\log}/S^{\log}} \rightarrow \mathrm{Ad}(P)$ with ψ' might not vanish. This composite will, however, vanish modulo I . Thus, it defines a horizontal morphism $\Phi_X^*(\tau_{X^{\log}/S^{\log}})_0 \rightarrow \mathcal{O}_{X_0}$, i.e., an element

$$\xi \in H^0(X_0, (\omega_{X^{\log}/S^{\log}})_0(-M_{X_0}^{\mathrm{ac}}))^F$$

(where the “ $-M_{X_0}^{\mathrm{ac}}$ ” follows from the fact that at the active marked points, the p -curvature is always 0 – cf. the discussion of Chapter II, §1.2). Note that ξ thus forms a section of a line bundle of *negative degree* ($= \deg(\omega_{Z^{\log}/S^{\log}}) - 2 = 1 - 2 = -1$) over Z_0 , so $\xi|_{Z_0} = 0$. Thus, we may regard ξ as an element of

$$H^0(Y_0, (\omega_{Y^{\log}/S^{\log}})_0(-M_{Y_0}^{\text{ac}}))^F$$

In fact, there is an ambiguity in the choice of κ' , so, taking this ambiguity into account, we obtain an element δ in the upper right-hand vector space of the commutative diagram at the end of §2.2. On the other hand, by the discussion at the end of §2.2, we see that by modifying our deformation by an appropriate element of $H_{\text{DR}}^1(X_0^{\log}, \text{Ad}(P_0))$ (which arises from an element of $H_{\text{DR}}^1(X_0^{\log}, \text{Ker}(\kappa_0))$ – note that $\text{Ker}(\kappa_0)$ is “ \mathcal{H} ” in the notation of §2.2), we may arrange so that $\delta = 0$. Finally, by adjusting κ' by an appropriate element of $H^0(X_0, \mathcal{K}_{X_0}) = H^0(Y_0, \mathcal{K}_{Y_0})$, we may even arrange that $\xi = 0$.

Thus, in summary, we see that we have shown that we may always lift $\{(P, \nabla_P)_I, \kappa_I\}$ to some $\{(P, \nabla_P), \kappa\}$, where (P, ∇_P) is of radii ρ , and κ is compatible with the p -curvature. Moreover, I *claim* that it follows from Proposition 2.10 (of §2.2), together with the calculation in the proof of Theorem 3.9 of Chapter II, that:

The dimension of the space of such deformations is precisely $3g - 3 + r$.

Indeed, the obstructions to the sort of deformations that interest us lie precisely in \mathcal{I}_{X_0} (which, by Proposition 2.10, satisfies $\dim_k(\mathcal{I}_{X_0}) = \dim_k(\mathcal{I}_{Y_0}) + 1$) and

$$H^0(Y, \omega_{Y^{\log}/S^{\log}}(-M_Y^{\text{ac}}))^F/\mathcal{D}$$

which may be identified with the space “ \mathcal{F}_0 ” in the proof of Theorem 3.9 of Chapter II. Moreover, by the proof of Theorem 3.9 of Chapter II, we have

$$2(3g - 3 + r - 1) - \dim_k(\mathcal{I}_{Y_0}) - \dim_k(\mathcal{F}_0) = 3g - 3 + r - 1$$

Thus, the dimension of the space of deformations that concern us here is equal to

$$\begin{aligned} 2(3g - 3 + r) - \dim_k(\mathcal{I}_{X_0}) - \dim_k(\mathcal{F}_0) &= 2 + \{2(3g - 3 + r - 1) - \dim_k(\mathcal{I}_{Y_0}) \\ &\quad - \dim_k(\mathcal{F}_0)\} - 1 \\ &= 1 + (3g - 3 + r - 1) \\ &= 3g - 3 + r \end{aligned}$$

as desired.

Thus, once we fix one of the two virtual p -curvatures (which denoted $\text{Ad}(P) \rightarrow \mathcal{V}_\beta$, $\text{Ad}(P) \rightarrow \mathcal{V}_\gamma$ in §2.2), we obtain a formal neighborhood of a k -valued point in $\overline{\mathcal{N}}_{g,r}^\rho$ which is smooth of dimension $3g - 3 + r$. Let us denote the corresponding formal neighborhoods by $N_\beta, N_\gamma \subseteq \overline{\mathcal{N}}_{g,r}^\rho$. Moreover, (from degree and monodromy considerations) if the restriction of $(P, \nabla_P)_0$ to Y_0^{\log} is mildly spiked of strength $d + a$, then it is clear that the generic element of N_β (respectively, N_γ) corresponds to a smooth, mildly spiked nilcurve of strength $d + b + p - c$ (respectively, $d + c + p - b$). The spiked locus of the p -curvature of this generic element is a divisor of order b (respectively, $p - b$) at the marked point that corresponds to the marked point β of Z_0^{\log} ; of order $p - c$ (respectively, c) at the marked point that corresponds to the marked point γ of Z_0^{\log} ; and the rest of this divisor forms a divisor which is flat over N_β (respectively, N_γ) of degree d and whose specialization to X_0^{\log} is contained in Y_0 . In particular, note that it follows that N_β and N_γ are distinct (since, for instance, $b \neq p - b$).

In summary, we have proven the following:

Theorem 2.13. *Let $X^{\log} \rightarrow S^{\log}$ be an r -pointed stable log-curve of genus g , where $S = \text{Spec}(k)$, and k is a field. Suppose that X^{\log} is obtained by gluing together an $(r - 1)$ -pointed stable log-curve Y^{\log} of genus g to a 3-pointed smooth log-curve Z^{\log} of genus 0 at a single node. Suppose further that X^{\log} is equipped with a nilpotent indigenous bundle (P, ∇_P) of radii ρ (in \mathbb{F}_p) whose restriction to Y^{\log} is mildly spiked of strength $d + a$, with a spike of order a at the node joining Y^{\log} to Z^{\log} ; while its restriction to Z^{\log} is dormant, with diameters b, c at the marked points of X^{\log} on Z . (Here we assume that $a + b + c$ is odd.) Let \mathcal{N} be the completion of $\overline{\mathcal{N}}_{g,r}^\rho$ at the k -valued point defined by $(X^{\log}, (P, \nabla_P))$.*

Then there exists a k -formally smooth closed subscheme of \mathcal{N} of dimension $3g - 3 + r$ whose generic point corresponds to a mildly spiked nilcurve of strength $d + b + p - c$ (respectively, $d + c + p - b$) and whose spikes at those marked points that correspond to the marked points of X^{\log} on Z are of orders b and $p - c$ (respectively, orders $p - b$ and c).

Finally, we consider what happens when the restriction of (P, ∇_P) to Y_0^{\log} is torally ordinary. Then I claim that

$$N_\beta \text{ is (formally) étale over } (\overline{\mathcal{M}}_{g,r})_{\mathbb{F}_p}.$$

Indeed, by the remarks following Definition 2.2, this is no problem with respect to the $3g - 3 + r - 1$ coordinates arising from deformations of Y_0^{\log} (i.e., deformations in which the critical point remains a node). Next, note that the formal completion of $\overline{\mathcal{N}}_{g,r}^\rho$ at $\{X_0^{\log}, (P, \nabla_P)_0\}$ is finite and flat of degree 2 over the formal completion of $(\overline{\mathcal{M}}_{g,r})_{\mathbb{F}_p}$ at X_0^{\log} (cf.

the discussion preceding Proposition 1.8 in Chapter II, §1.3). Thus, the claim follows from the following lemma in commutative algebra:

Lemma 2.14. *Let $A = k[[t_1, \dots, t_n]]$ (where t_1, \dots, t_n are indeterminates). Let B be a finite, flat A -algebra which is of rank 2 as an A -module. Suppose further that B admits two (ring) quotients $B \rightarrow N_1; B \rightarrow N_2$, where N_1 and N_2 are formally smooth of dimension n over k . Assume further that neither of these two quotients factors through the other, even after tensoring with the quotient field $Q(A)$ of A . Then it follows that $N_1 \cong N_2 \cong A$ as A -algebras.*

Proof. Indeed, N_1 is a finite A -algebra, and both A and N_1 are regular rings of the same dimension, so N_1 is faithfully flat over A . Similarly, N_2 is faithfully flat over A . Since $\dim_{Q(A)}(B \otimes Q(A)) = 2$, and neither of $B \otimes Q(A) \rightarrow N_1 \otimes Q(A)$ and $B \otimes Q(A) \rightarrow N_2 \otimes Q(A)$ factors through the other, it follows that $\dim_{Q(A)}(N_1 \otimes Q(A)) = \dim_{Q(A)}(N_2 \otimes Q(A)) = 1$. This, combined with faithful flatness, implies that $N_1 = N_2 \cong A$, as desired.

○

Thus, we have proven the following result:

Corollary 2.15. *Suppose that the restriction of $(P, \nabla_P)_0$ to Y_0^{\log} is torally ordinary. Then both N_β and N_γ map isomorphically to the completion of $(\overline{\mathcal{M}}_{g,r})_{\mathbb{F}_p}$ at the k -valued point of $(\overline{\mathcal{M}}_{g,r})_{\mathbb{F}_p}$ defined by X_0^{\log} . Moreover, the nilcurves corresponding to the generic points of $N_\beta, N_\gamma \subseteq \overline{\mathcal{N}}_{g,r}^\rho$ are torally ordinary.*

Remark. Note that in the aphillal case, we derived the main result concerning how the spiked locus degenerates (i.e., the last part of Theorem 2.5) by means of *degeneration theory over a trait*, whereas in the “grafting on a dormant atom case,” we derived the main result concerning how the spiked locus degenerates (i.e., Theorem 2.13) by means of *deformation theory*. It is probably possible to handle the aphillal case by means of deformation theory, or the “grafting on a dormant atom case” by means of degeneration theory, but I believe that the approaches chosen here are easier and more natural. Also, in the “grafting on a dormant atom case,” it is not clear how to show that *both* possible virtual p -curvatures actually occur (and with “multiplicity 1”) if one uses degeneration theory. On the other hand, deformation theory (being, after all, somewhat more precise) is typically technically much more difficult to set up than degeneration theory (cf. the fact that §2.2 and 2.3 together are longer than §2.1), while in the aphillal case (unlike the dormant atom case) we only have one possible outcome for the generic p -curvature, so there is no need to show that “all possible outcomes actually occur, and with the correct multiplicities.”

§2.4. Proof of the Main Theorem

We begin with an elementary but important Lemma from commutative algebra:

Lemma 2.16. *Let (A, \mathfrak{m}_A) be an artinian local ring which is a local complete intersection. Then the length of A is $\geq 2^r$, where $r \stackrel{\text{def}}{=} \dim_K(\mathfrak{m}_A/\mathfrak{m}_A^2)$, and K is the residue field of A . (Recall from §1 that when equality holds, we refer to A as “taut.”)*

Proof. By definition, there exists a complete regular local ring B together with a surjection $B \rightarrow A$ whose kernel I is generated by $\dim(B)$ elements. It is easy to see that without loss of generality, we may assume that $\dim(B) = r$. Then it follows that $I \subseteq \mathfrak{m}_B^2$ (where $\mathfrak{m}_B \subseteq B$ is the maximal ideal of B). Let us write $I = (f_1, \dots, f_r)$. Let $I_i \stackrel{\text{def}}{=} (f_1, \dots, f_i)$; $B_i \stackrel{\text{def}}{=} B/I_i$ (for $i = 1, \dots, r$). Moreover, we let $B_0 \stackrel{\text{def}}{=} B$. Let $P_i(X)$ be the unique polynomial in the indeterminate X such that for n sufficiently large, we have $P_i(n) = \ln(B_i/\mathfrak{m}_{B_i}^n)$ (where “ \ln ” stands for “length,” and $\mathfrak{m}_{B_i} \subseteq B_i$ is the maximal ideal of B_i). Then one knows (see, e.g., [Mats]) that $P_i(X)$ is a polynomial of degree $r-i$, with leading term of the form $\frac{d_i}{(r-i)!} X^{r-i}$ (for some nonzero $d_i \in \mathbf{Q}$). We would like to prove (by induction on i) the following statement:

(*) We have $d_i \geq 2^i$.

Note that since $d_r = P_r(X) = \ln(B_r) = \ln(A)$, this will complete the proof of the Lemma.

Note that $d_0 = 1$. Now assume that $d_i \geq 2^i$, for some $i \in \{0, \dots, r-1\}$. Note that multiplication by f_{i+1} in B_i is injective, hence induces an exact sequence

$$0 \rightarrow B_i \rightarrow B_i \rightarrow B_{i+1} \rightarrow 0$$

Tensoring this exact sequence with B/\mathfrak{m}_B^n then gives an exact sequence

$$B_i/\mathfrak{m}_{B_i}^n \rightarrow B_i/\mathfrak{m}_{B_i}^n \rightarrow B_{i+1}/\mathfrak{m}_{B_{i+1}}^n \rightarrow 0$$

On the other hand, since $f_{i+1} \in \mathfrak{m}_B^2$, it follows that the first arrow in this last exact sequence factors through $B_i/\mathfrak{m}_{B_i}^{n-2}$. Thus, applying the length function, we see that we obtain (for n large) $P_{i+1}(n) \geq P_i(n) - P_i(n-2)$. Moreover, $P_i(X) - P_i(X-2)$ is a polynomial of degree $r-i-1$ whose leading term is equal to $\frac{2d_i}{(r-i-1)!} X^{r-i-1}$. Thus, by comparing

leading terms, we obtain $d_{i+1} \geq 2d_i$. This completes the induction step in the proof of (*), and hence also the proof of the Lemma. \circ

Now we would like to move on to the proof of Theorem 1.1. First, let us take care of the *dormant case* (i.e., (1) in the statement of Theorem 1.1), since it is the easiest case, and, moreover, demonstrates the essential idea involved. First of all, we recall that it was proven in Theorem 2.8 of Chapter II that there exists a closed substack $\mathcal{N}' \subseteq \mathcal{M}$ which is étale over (in fact, isomorphic to) \mathcal{M} and whose generic point corresponds to a dormant nilcurve. The main issues then are to determine the padding degree of $\mathcal{N}'_{\bar{\eta}}$ at the generic point of \mathcal{N}' , and to show that $\mathcal{N}' = \mathcal{N}_{\text{red}}$ (i.e., to show that \mathcal{N}_{red} does not contain any irreducible components other than \mathcal{N}'). On the other hand, since this generic point parametrizes a *dormant* nilcurve, it is clear that the morphism Φ_P of Definition 2.2 vanishes identically for this nilcurve. Thus, it follows that the Zariski cotangent space at this generic point is of dimension $3g - 3 + r$. But then Lemma 2.16 implies that the padding degree of $\mathcal{N}'_{\bar{\eta}}$ at this generic point is $\geq 2^{3g-3+r}$. Moreover, the entire degree of $\mathcal{N}'_{\bar{\eta}}$ over $\bar{\eta}$ is 2^{3g-3+r} . Thus, it follows that the padding degree is precisely 2^{3g-3+r} , and, moreover, that $\mathcal{N}_{\text{red}} = \mathcal{N}'$, as desired.

Now we move on to the more complicated *nondormant case* (i.e., (2) of Theorem 1.1). We begin by stating that we shall use induction on $3g - 3 + r$. Thus, we shall, in the following discussion, always be free to assume that Theorem 1.1 is known to be true for smaller $3g - 3 + r$. Note, in connection with this induction, that by a well-known descent argument (plus the fact that $(\overline{\mathcal{M}}_{g,r})_{\mathbb{F}_p}$ is of finite type over \mathbb{F}_p , hence *excellent*), the various statements of Theorem 1.1 hold as stated in Theorem 1.1 if and only if their corresponding “strict henselian versions” hold. That is to say, if we let \mathcal{M}^{sh} (respectively, \mathcal{N}^{sh}) be the strict henselization of $(\overline{\mathcal{M}}_{g,r})_k$ (respectively, $(\overline{\mathcal{N}}_{g,r})_k$) at the point defined by the underlying curve of the molecule M (respectively, point defined by M), and $\bar{\eta}^{\text{sh}}$ be the strict henselization of the generic point of \mathcal{M}^{sh} , then Theorem 1.1 holds as stated if and only if it holds *mutatis mutandis* with \mathcal{N} , \mathcal{M} , and $\bar{\eta}$ replaced by \mathcal{N}^{sh} , \mathcal{M}^{sh} , and $\bar{\eta}^{\text{sh}}$. The reason we make remark is that when applying the induction step (i.e., deforming once, restricting to a generic point of this deformation, then deforming again in a neighborhood of this generic point, etc.), it is often useful to replace completions by strict henselizations.

Next, let us fix a *plot* Π for M . Then for each scenario Σ of (M, Π) , we would like to define an \mathcal{N}_{Σ} as in the statement of the Theorem. To do this, let us first choose an ordering Π_{Σ} of *all* the nodes of M as follows: The first nodes in this ordering are the nodes of ν_{Π} in the order given by Π . Next, we place the aphial and classical ordinary nodes (in any order). Finally, last in the ordering Π , we have the phial nodes (in any order).

The point now (as should be obvious from the way the various

terms were defined in §1) is that one deforms, one after the other, each of the nodes on M in the order specified by Π_Σ . Here, the deformation theories that we use are:

- (1) §2.3 of the present Chapter for *grafted nodes*;
- (2) §3 of Chapter II for *classical ordinary nodes*;
- (3) §2.1 of the present Chapter for *aphilial nodes* (here, we regard Theorem 2.5 of §2.1 as a result in “deformation theory” in the sense that it says that at an aphilial node, the spiked locus “necessarily deforms to the empty set”); and
- (4) §3 of Chapter II for *philial nodes*.

When we graft on a dormant atom (i.e., in the terminology of §1, deform at a grafted node), then according to the theory of §2.3, we must choose one of the two possible virtual p -curvatures. It is this choice which is given by the *scenario* Σ that one uses. Thus, each choice in the scenario entails the selection of one of the two distinct “sheets” (i.e., N_β and N_γ in the notation of §2.3). Moreover, the signs σ_x of Σ at the niches corresponding to marked points x are such that $\sigma_x(2\rho_x)$ is precisely the vanishing order at x of the p -curvature of the deformed nilcurve. This is the reason for the condition (in the definition of a “scenario” – see §1) that the product of the signs of the three niches of a newly adjoined dormant atom be “–.” That is to say, this condition corresponds to the condition imposed at the beginning of §2.2 that the parity of $a+b+c$, which is the same as the parity of $(p-a)+b+(p-c)$ or the parity of $(p-a)+(p-b)+c$ (cf. the last paragraph of the statement of Theorem 2.13), be odd.

At any rate, deforming each of the nodes on M in the order specified by Π_Σ gives rise to a reduced $\mathcal{N}'_\Sigma \subseteq \mathcal{N}_{\bar{\eta}}$. Note that *the only sort of deformation that can give rise to a new spike on the resulting smooth nilcurve is a deformation at a philial node*. This shows that the resulting smooth nilcurve parametrized by \mathcal{N}'_Σ is necessarily mildly spiked of the desired strength, $\Sigma(\rho) + p \cdot n_{\text{phl}}(\Sigma)$. Let \mathcal{N}_Σ be the open substack of $\mathcal{N}_{\bar{\eta}}$ with the same underlying topological space as \mathcal{N}'_Σ .

Next, we *claim* that the degree of \mathcal{N}'_Σ over $\bar{\eta}$ is $\geq 2^{n_{\text{aph}}(\Sigma)}$. Indeed, by the induction hypothesis, it suffices to handle the case where *there are no philial nodes*. On the other hand, if there are no philial nodes, then Corollaries 2.6 and 2.15 imply that the smooth nilcurve parametrized by \mathcal{N}'_Σ is *totally ordinary*, hence that \mathcal{N}_Σ is étale over $\bar{\eta}$ at \mathcal{N}'_Σ ; thus, $\mathcal{N}_\Sigma = \mathcal{N}'_\Sigma$. But deformation at an aphilial node always adds a factor of two to the degree (over $\bar{\eta}$) of \mathcal{N}_Σ (cf. Chapter II, §1.2, 1.3), so the claim that $\mathcal{N}'_\Sigma \geq 2^{n_{\text{aph}}(\Sigma)}$ thus follows from the induction hypothesis.

Thus, the main issues that remain are determining the various degrees (over $\bar{\eta}$) and padding degrees, and showing that the \mathcal{N}'_Σ (as

Σ ranges over all $2^{n_{\text{dor}}}$ possible scenarios) exhaust $(\mathcal{N}_{\bar{\eta}})_{\text{red}}$. (Note that it is already clear that the various \mathcal{N}'_{Σ} (corresponding to distinct Σ) do not intersect one another.) Just as in the dormant case, these two issues are related and can be settled by using Lemma 2.16 as follows:

First, let M'_{Σ} be the nilcurve parametrized by \mathcal{N}'_{Σ} . Let X_{Σ}^{\log} be the underlying (smooth) curve of M'_{Σ} . The portion of the vanishing locus of the p -curvature of M'_{Σ} that avoids the marked points is the Frobenius pull-back of a divisor

$$D_{\Sigma} \subseteq X_{\Sigma}^F$$

of degree $n_{\text{phl}}(\Sigma)$ in X_{Σ}^F . Next, observe that the morphism for M'_{Σ} which was called “ Ψ_P ” in the discussion preceding Definition 2.2 factors through the submodule

$$H^0(X_{\Sigma}^F, (\omega_{X_{\Sigma}^{\log}/\mathcal{N}'_{\Sigma}}^{\otimes 2}(-M_{X_{\Sigma}^{\log}}))^F(-D_{\Sigma})) \subseteq H^0(X_{\Sigma}, (\omega_{X_{\Sigma}^{\log}/\mathcal{N}'_{\Sigma}}^{\otimes 2}(-M_{X_{\Sigma}^{\log}})))^F$$

– indeed, this follows immediately from the fact that “ Ψ_P ” was defined as the composite of the (dual of the) p -curvature with another morphism (plus the fact that the p -curvature vanishes on the Frobenius pull-back of D_{Σ}). Now we have the following result:

Lemma 2.17. *This submodule*

$$H^0(X_{\Sigma}^F, (\omega_{X_{\Sigma}^{\log}/\mathcal{N}'_{\Sigma}}^{\otimes 2}(-M_{X_{\Sigma}^{\log}}))^F(-D_{\Sigma})) \subseteq H^0(X_{\Sigma}, (\omega_{X_{\Sigma}^{\log}/\mathcal{N}'_{\Sigma}}^{\otimes 2}(-M_{X_{\Sigma}^{\log}})))^F$$

is free of rank $3g - 3 + r - n_{\text{phl}}(\Sigma)$ over $\mathcal{O}_{\mathcal{N}'_{\Sigma}}$.

Proof. We apply Riemann-Roch on the curve X_{Σ}^F . Thus, it suffices to show that

$$H^0(X_{\Sigma}^F, \omega_{(X_{\Sigma}^{\log})^F/\mathcal{N}'_{\Sigma}}^{-1})(D_{\Sigma})) = 0$$

But this will follow if we can show that the line bundle

$$\mathcal{G} \stackrel{\text{def}}{=} \omega_{(X_{\Sigma}^{\log})^F/\mathcal{N}'_{\Sigma}}(-D_{\Sigma})$$

is ample on X_{Σ}^F . On the other hand, if P_{Σ} is the \mathbf{P}^1 -bundle on X_{Σ} (underlying the nilcurve M'_{Σ}), then the dual of the p -curvature of this nilcurve defines a horizontal morphism

$$\mathrm{Ad}(P_\Sigma) \rightarrow \Phi_{X_\Sigma}^*(\mathcal{G})$$

Now the fact that $(P_\Sigma, \nabla_{P_\Sigma})$ is *torally crys-stable* (Chapter I, Definition 1.2, (3)) implies that \mathcal{G} is ample. This completes the proof. \circ

Thus, Lemma 2.17 implies that the rank of the morphism “ Ψ_P ” associated to M'_Σ is $\leq 3g - 3 + r - n_{\mathrm{phl}}(\Sigma)$. In particular, the kernel of Φ_P will have rank $\geq 3g - 3 + r + n_{\mathrm{phl}}(\Sigma)$. On the other hand, this kernel is precisely the dimension of the Zariski cotangent space to $\overline{\mathcal{N}}_{g,r}$ at a general closed point of $\overline{\mathcal{N}}_{g,r}$ lying in the irreducible components of $\overline{\mathcal{N}}_{g,r}$ defined by \mathcal{N}'_Σ . Since $\overline{\mathcal{N}}_{g,r}$ is $(3g - 3 + r)$ -dimensional, this implies that the dimension of the Zariski cotangent space of \mathcal{N}_Σ (at every point of \mathcal{N}_Σ) is $\geq n_{\mathrm{phl}}(\Sigma)$. Thus, by Lemma 2.16, we obtain that *the padding degree of \mathcal{N}_Σ at each of its points is $\geq 2^{n_{\mathrm{phl}}(\Sigma)}$* .

Now we put everything together. Note that the various observations made so far imply that the degree of \mathcal{N}_Σ over $\overline{\eta}$ is

$$\geq 2^{n_{\mathrm{aph}}(\Sigma) + n_{\mathrm{phl}}(\Sigma)} = 2^{n_{\mathrm{tor}} - n_{\mathrm{dor}}}$$

Adding up over all $2^{n_{\mathrm{dor}}}$ scenarios Σ , we thus obtain that the union of all the \mathcal{N}_Σ 's has degree $\geq 2^{n_{\mathrm{tor}}}$ over $\overline{\eta}$. But $2^{n_{\mathrm{tor}}}$ is precisely the degree of \mathcal{N} over \mathcal{M} . This shows simultaneously (as in the dormant case) that the \mathcal{N}_Σ exhaust \mathcal{N} , and that the degrees and padding degrees of the \mathcal{N}_Σ are as desired. The rest of Theorem 1.1 then follows formally from everything that we have done so far. *This completes the proof of Theorem 1.1.*

§3. Examples

In this §, we work out various consequences and special cases of the main theorem of the Chapter.

§3.1. Consequences in the Case $(g,r)=(1,1)$

In this subsection, we observe that the theory of §1 (the aphillial case) allows us to compute the degree of the *nonordinary locus* of $\mathcal{N}_{1,1}$. Indeed, let $\overline{\mathcal{N}} \stackrel{\mathrm{def}}{=} \overline{\mathcal{N}}_{1,1}$; $\mathcal{N} \stackrel{\mathrm{def}}{=} \mathcal{N}_{1,1}$; $\overline{\mathcal{M}} \stackrel{\mathrm{def}}{=} \overline{\mathcal{M}}_{1,1}$; $\mathcal{M} \stackrel{\mathrm{def}}{=} \mathcal{M}_{1,1}$; $\overline{\mathcal{S}} \stackrel{\mathrm{def}}{=} \overline{\mathcal{S}}_{1,1}$ (the Schwarz torsor over $\overline{\mathcal{M}}$); $\overline{\mathcal{Q}} \stackrel{\mathrm{def}}{=} \overline{\mathcal{Q}}_{1,1}$ (as in the discussion of the Verschiebung morphism in Chapter II, §1.3). Let λ be the Hodge bundle on $\overline{\mathcal{M}}$, and $\infty \subseteq \overline{\mathcal{M}}$ be the divisor at infinity. Let $\Omega \stackrel{\mathrm{def}}{=} \lambda^{\otimes 2}$.

Thus, $\Omega = \Omega_{\overline{\mathcal{M}}/\mathbb{F}_p}(\infty)$. The morphism “ Φ_P ” of Definition 2.2 defines a morphism $\Omega|_{\overline{\mathcal{N}}} \rightarrow \Omega^{\otimes p}|_{\overline{\mathcal{N}}}$, whose zero locus is some divisor $D_{\overline{\mathcal{N}}} \subseteq \overline{\mathcal{N}}$. Write $D_{\mathcal{N}}$ for the restriction of $D_{\overline{\mathcal{N}}}$ to \mathcal{N} .

Definition 3.1. We define the *nonordinary locus* of $\mathcal{N} = \mathcal{N}_{1,1}$ to be the divisor $D_{\mathcal{N}}$.

In the case of stacks, there is some question how to define the notion of “degree.” Here we define it by means of the normalization $\deg(\infty) = 1$. Thus, in order to compute the degree of $D_{\mathcal{N}}$, it suffices to compute $\deg(D_{\overline{\mathcal{N}}})$, and the degree of $D_{\overline{\mathcal{N}}}$ at infinity. Now one knows that $\deg(\lambda) = \frac{1}{12}$, so it follows that

$$\deg(D_{\overline{\mathcal{N}}}) = (p-1) \cdot \deg(\Omega|_{\overline{\mathcal{N}}}) = p(p-1) \cdot \deg(\Omega) = \frac{p(p-1)}{6}$$

On the other hand, it follows from the proof of Corollary 2.6 that the zero divisor of the morphism $\Omega|_{\overline{\mathcal{N}}} \rightarrow \Omega^{\otimes p}|_{\overline{\mathcal{N}}}$ used to define $D_{\overline{\mathcal{N}}}$ is equal to $a \cdot (\frac{1}{2} \cdot \infty)$ at the point of $\overline{\mathcal{N}}$ corresponding to the molecule whose diameter at its unique node is equal to a (where $a = 0, \dots, \frac{p-1}{2}$). Since the degree of the completion of $\overline{\mathcal{N}}$ at this point (over the corresponding completion of $\overline{\mathcal{M}}$) is 2, it thus follows that the degree (as a divisor on $\overline{\mathcal{N}}$) of the portion of $D_{\overline{\mathcal{N}}}$ which is supported on this completion is equal to a . Thus, adding up all these degrees (i.e., all these “ a ”’s), we obtain

$$\frac{1}{2} \left(\frac{p-1}{2} \right) \left(\frac{p+1}{2} \right) = \frac{p^2-1}{8}$$

Subtracting this from $\deg(D_{\overline{\mathcal{N}}})$, we thus obtain:

Corollary 3.2. The degree of the nonordinary locus of $\mathcal{N}_{1,1}$ is equal to $\frac{(p-1)(p-3)}{24}$ (where degree is defined so that for $\infty \in \overline{\mathcal{M}}_{1,1}$, we have $\deg(\infty) = 1$).

We remark that in the case $p = 5$, the degree obtained in Corollary 3.2 is compatible with the results of Chapter IV, §1.3 (cf. especially the second Remark following Theorem 1.4 of Chapter IV, §1.3). Indeed, in this case, \mathcal{N} consists of two connected components, one which is étale of degree 3 over \mathcal{M} , and one which is of degree 2 over \mathcal{M} and ramified precisely at the point corresponding to an elliptic curve with an automorphism of order three.

Finally, let us observe that $D_{\overline{\mathcal{N}}}$ may be identified with the ramification divisor of $\overline{\mathcal{N}}$ over $\overline{\mathcal{M}}$. Indeed, recall that $\overline{\mathcal{N}} \subseteq \overline{\mathcal{S}}$ is defined as the pull-back of the zero section of $\overline{\mathcal{Q}} \rightarrow \overline{\mathcal{M}}$ via the *Verschiebung morphism* (cf. Chapter II, §1.3):

$$\overline{\mathcal{S}} \rightarrow \overline{\mathcal{Q}}$$

Moreover, Ω (respectively, $\Omega^{\otimes p}$) may be identified with the relative tangent bundle of $\overline{\mathcal{S}}$ (respectively, $\overline{\mathcal{Q}}$) over $\overline{\mathcal{M}}$. Finally, it is easy to check that the morphism $\Omega|_{\overline{\mathcal{N}}} \rightarrow \Omega^{\otimes p}|_{\overline{\mathcal{N}}}$ (i.e., the morphism “ Φ_P ” of Definition 2.2) may be identified with the map on relative tangent bundles obtained by differentiating the Verschiebung morphism at the point $\overline{\mathcal{N}} \subseteq \overline{\mathcal{S}}$. This completes the verification that $D_{\overline{\mathcal{N}}}$ may be identified with the ramification divisor of $\overline{\mathcal{N}}$ over $\overline{\mathcal{M}}$. In particular, this allows us to compute the local structure of $\overline{\mathcal{N}}$ at infinity:

Corollary 3.3. *For $a = 0, \dots, \frac{p-1}{2}$, the structure (as an $\hat{\mathcal{O}}_{\overline{\mathcal{M}}, \infty}$ -algebra) of the completion $\hat{\mathcal{O}}_{\overline{\mathcal{N}}, \infty_a}$ of the local ring of $\overline{\mathcal{N}}$ at the point ∞_a of $\overline{\mathcal{N}}$ corresponding to the molecule whose diameter at its unique node is equal to a is given by: $\hat{\mathcal{O}}_{\overline{\mathcal{N}}, \infty_a} \cong \hat{\mathcal{O}}_{\overline{\mathcal{M}}, \infty}[x]/(x^2 - \pi^a)$, where $\pi \in \hat{\mathcal{O}}_{\overline{\mathcal{M}}, \infty}$ is a generator of the maximal ideal of $\hat{\mathcal{O}}_{\overline{\mathcal{M}}, \infty}$.*

Results (whose statements and proofs are entirely similar) analogous to Corollaries 3.2 and 3.3 can also be obtained for $(g, r) = (0, 4)$. We leave it to the reader to make these explicit.

§3.2. Explicit Computations

In this subsection, we use Theorem 1.1 to perform the calculations summarized in Corollary 1.3. First observe that the statements in Corollary 1.3, (1), (2) concerning *continuous integrals* follow immediately from the corresponding statements concerning discrete integrals, plus elementary calculus (the theory of Riemann sums). Moreover, to obtain the discrete parts of Corollary 1.3, (1), (2), we simply apply Theorem 1.1 to the following totally degenerate curves: when $g = 0$, we use the curve which is a chain of \mathbf{P}^1 's; when $g = 1$, we use the curve which is a ring of \mathbf{P}^1 's. Next, we turn to part (3) of Corollary 1.3. The cases $(0, 4)$; $(1, 1)$; and $(1, 2)$ are clear (since in these cases, all smooth nilcurves are automatically admissible). To complete the rest of the calculations, the following elementary lemmas will be useful:

Lemma 3.4. *We have:*

$$\sum_{\alpha=1}^r \alpha = \frac{1}{2}r(r+1); \quad \sum_{\alpha=1}^r \alpha^2 = \frac{1}{6}r(r+1)(2r+1)$$

where $r \geq 1$ is an integer.

Proof. One simply uses induction on r . \circ

Lemma 3.5. *Let n be a nonnegative integer. Let \mathcal{A}_n be the set of triples (a, b, c) of positive integers a, b, c such that $a + b, a + c, b + c \leq n + 1$. Then the cardinality $|\mathcal{A}_n|$ of \mathcal{A}_n is given by:*

$$\frac{n}{8}(2n^2 + 3n + 2) + \frac{\delta_n}{8}$$

where $\delta_n = 0$ (respectively, $\delta_n = 1$) if n is even (respectively, odd).

Proof. Let $h(n) \stackrel{\text{def}}{=} |\mathcal{A}_n| - |\mathcal{A}_{n-1}|$. Then $h(n)$ can be computed easily using Lemma 3.4: If m is a positive integer, then $h(2m) = 3m^2$; $h(2m + 1) = 3m(m + 1) + 1$. Now adding up h 's via Lemma 3.4 gives the desired formula. \circ

Lemma 3.6. *Let n be a positive even integer. Let \mathcal{B}_n be the set of triples (a, b, c) of positive integers a, b, c such that $a + b + c$ is even and $\leq n$. Then $|\mathcal{B}_n|$ is given by:*

$$\frac{1}{6}(4m^3 - 3m^2 - m)$$

where we write $n = 2m$.

Proof. Let $k(n) \stackrel{\text{def}}{=} |\mathcal{B}_n| - |\mathcal{B}_{n-2}|$. By Lemma 3.4, we compute that $k(n) = \frac{1}{2}(n - 1)(n - 2)$. Then adding up (using Lemma 3.4) gives the result. \circ

It follows easily from the results on atoms summarized at the beginning of §1 that $|\mathcal{B}_{(p-1)}|$ is equal to the number of (isomorphism classes of) nondormant atoms (over \mathbb{F}_p) all of whose radii are nonzero. Thus, we obtain that:

Corollary 3.7. *In characteristic p , the number of isomorphism classes of dormant atoms is equal to $\frac{1}{24}p(p^2 - 1)$.*

Proof. This follows by subtracting $|\mathcal{B}_{(p-1)}|$ from $(\frac{p-1}{2})^3$ (= the number of isomorphism classes of atoms all of whose radii are nonzero). \circ

Now let us consider $n_{2,0}^{\text{ord}}$. Since every generic nilcurve of type $(2,0)$ is either dormant or ordinary, it thus follows that $n_{2,0}^{\text{ord}}$ is obtained by subtracting from p^3 (the degree of $\mathcal{N}_{2,0}$ over $(\mathcal{M}_{2,0})_{\mathbb{F}_p}$) 8 times the number obtained in Corollary 3.7. (Here the 8 arises from the “padding degree.”) This gives the expression for $n_{2,0}^{\text{ord}}$ stated in Corollary 1.3, (3).

Next, let us consider the case $(g,r) = (0,5)$. Plugging into the formula of Corollary 1.3, (1), we obtain (by Lemma 3.4)

$$n_{0,5}^{\text{ord}} = 4 \cdot \left(\frac{p-1}{8}\right)(p-3) + 2 \cdot (p-1) + 1$$

(Here, since we are computing the first “integral” in Corollary 1.3, (1), we group the terms according to the value of “ N_α .” That is to say, the first term (beginning with a “4.”) corresponds to the case of precisely two nonzero radii; the second term (beginning with a “2.”) corresponds to the case of precisely one nonzero radius; and the last term (“1”) corresponds to the case where all radii are zero. We shall continue to use this convention throughout the following discussion.) Simplifying this expression gives the expression stated in Corollary 1.3, (3). One can check this expression by computing (by means of Theorem 1.1) the degree of the nonordinary portion of $\mathcal{N}_{0,5}$ over $(\mathcal{M}_{0,5})_{\mathbb{F}_p}$, and showing that these two degrees add up to p^2 .

Next, we consider the case $(g,r) = (0,6)$. Plugging into the formula of Corollary 1.3, (1), we obtain (by Lemma 3.4)

$$8 \cdot \left\{ \frac{1}{24}(p-1)(p-2)(p-3) \right\} + 4 \cdot \left\{ \left(\frac{p-1}{2}\right)^2 + \frac{1}{4}(p-1)(p-3) \right\} + 2 \cdot \left\{ \frac{3}{2}(p-1) \right\} + 1$$

Simplifying this expression gives the expression stated in Corollary 1.3, (3). One can check this expression by computing $n_{0,6}^{\text{ord}}$ by means of the totally degenerate curve obtained by taking a copy X^{\log} of $(\mathbb{P}^1)^{\log}$ (with three marked points), and gluing three more copies of $(\mathbb{P}^1)^{\log}$ to X^{\log} at the three marked points of X^{\log} . We shall refer to the irreducible component X^{\log} as “central.” Thus, the six marked points all lie on the three non-central copies of $(\mathbb{P}^1)^{\log}$ (i.e., two marked points on each). Now, by using Corollary 3.7 for the first expression in brackets – which corresponds to the case where the central atom is *dormant* – and using Lemma 3.4 for the summations in the other terms – which correspond to the various cases where the central atom is *nondormant*, we obtain:

$$4 \cdot \left\{ \frac{1}{24}p(p^2-1) \right\} + 8 \cdot \left\{ \frac{1}{48}(p-1)(p-3)(p-5) \right\} + 4 \cdot \left\{ \frac{3}{8}(p-1)(p-3) \right\} + 2 \cdot \left\{ \frac{3}{2}(p-1) \right\} + 1$$

One checks that simplifying this expression again gives the expression stated in Corollary 1.3, (3).

Finally, we consider the case $(g, r) = (1, 3)$. Plugging into the formula of Corollary 1.3, (2), we obtain (by Lemmas 3.4 and 3.5):

$$8 \cdot \left\{ \frac{1}{16}(p^3 - 3p^2 + 7p - 5) \right\} + 4 \cdot \left\{ \frac{3}{8}(p-1)(p-3) \right\} + 2 \cdot \left\{ \frac{3}{2}(p-1) \right\} + 1$$

Simplifying this expression gives the expression stated in Corollary 1.3, (3). One can check this expression by computing $n_{1,3}^{\text{ord}}$ by means of the following totally degenerate curve: Let X^{log} be a totally degenerate 1-pointed curve of genus 1. Let Y^{log} be a totally degenerate 4-pointed curve of genus 0. Then the curve that we will use is the one obtained by gluing X^{log} (at its unique marked point) to Y^{log} (at some marked point of Y^{log}). If we first deform the 1-pointed curve of genus 1 into a smooth curve (which gives us a factor of “ p ” in each term below), and then compute the sums resulting from the remaining two nodes by means of Lemma 3.4, we obtain:

$$4 \cdot \left\{ \frac{p}{8}(p-1)(p-3) \right\} + 2 \cdot \{p(p-1)\} + p$$

which also agrees with the expression given in Corollary 1.3, (3). This completes the proof of Corollary 1.3.

Remark. Note that the second approach given above to computing $n_{1,3}^{\text{ord}}$ is technically much easier than the first approach (which uses Lemma 3.5, which is by far the least trivial of the various computational Lemmas given above). This observation suggests that perhaps many of the combinatorial identities inherent in Theorem 1.1 are (at least computationally) highly nontrivial.


Remark. One of the main reasons why I was particularly interested in computing $n_{0,6}^{\text{ord}}$ and $n_{1,3}^{\text{ord}}$ in two essentially different ways and seeing that the results agreed was to assure myself, by means of concrete evidence, of the correctness of the rather extensive and abstract theory behind Theorem 1.1.

Pictorial Appendix

In this Appendix, we provide illustrations of the five types of nodal deformation discussed in Theorem 1.1. The lightly shaded regions represent the nondormant portion of the nilcurve; the darkly shaded


regions represent the dormant portion of the nilcurve; the signs at the various niches are the signs associated to these niches by the scenario in question; and, finally, the dark spot in the phial case is a "spike," i.e., a punctual divisor of degree p where the p -curvature vanishes.

The Dormant Case:




The Nondormant Case:


A Classical Ordinary Node:




A Grafted Node:



A Phial Node:



An Aphial Node:



Chapter VI: The Stack of Quasi-Analytic Self-Isogenies

§0. Introduction

So far, in the first five chapters of this work, we have mainly been concerned with the characteristic p theory. In this Chapter, we introduce the main characters “ \mathcal{Q} ” in the theory of p -adic liftings.

The purpose of these main characters is to serve as moduli stacks of Frobenius invariant indigenous bundles.

More concretely: In Chapter III, we defined various “VF-stacks” which parametrize certain data in characteristic p . In this Chapter, we consider stacks “ \mathcal{Q} ” that parametrize Frobenius invariant *liftings* of this data to \mathbf{Z}_p . Since we want to deal with *Frobenius invariant* liftings, the first issue that arises is what sorts of schemes we will work over. That is to say, the stack “ \mathcal{Q} ” exists only in the sense that for any object S in some appropriate category, we assign a value (i.e., a groupoid) $\mathcal{Q}(S)$ to S . Thus, we must specify this “appropriate category.” On the one hand, we want this category to be big enough to contain various natural objects like Witt schemes and “ B_{crys} ”; on the other hand, the category must be small or special enough so that its objects admit natural Frobenius liftings – a very stringent condition – otherwise, we are unable to define a notion of “Frobenius invariant.” The solution to this problem is the category of *epiperfect schemes*. The first § of this Chapter is thus devoted to developing the theory of this category and defining the stack \mathcal{Q} .

Once one has defined \mathcal{Q} , one wants to *compactify* it.

The desire to compactify \mathcal{Q} is in some sense a desire to generalize the basic fact that $\overline{\mathcal{N}}_{g,r}$ is proper over \mathbf{F}_p .

That is to say, “ideally,” one might hope to “compactify” the ordinary theory of [Mzk1] by extending it to all of $\overline{\mathcal{N}}_{g,r}$, and thus obtaining a natural lifting of $\overline{\mathcal{N}}_{g,r}$ over \mathbf{Z}_p . Unfortunately, it is not difficult to see that this is unrealistic: in order to obtain a theory even remotely reminiscent of [Mzk1] (in the sense that one deals with Frobenius invariant indigenous bundles of some sort), it seems fairly clear that already one must restrict oneself to working over such stacks as the VF-stacks $\overline{\mathcal{N}}_{g,r}^\Pi$. However, as we saw in Chapter III, Theorem 2.10, these stacks are quasi-affine, hence far from proper. Nonetheless, we still wish to “compactify.” The compromise here is obtained by adjusting our notion of what it means to be “proper.” To do this, first we observe that all the \mathcal{Q} ’s obtained in §1 are naturally contained in $(\overline{\mathcal{S}}_{g,r})_W$ (where $\overline{\mathcal{S}}_{g,r}$ is the stack of stable curves equipped with an indigenous bundle, and “ W ” denotes a sort of “infinite Weil restriction of scalars” from \mathbf{Z}_p to \mathbf{F}_p). In other words, $(\overline{\mathcal{S}}_{g,r})_W$ (and its companion $(\overline{\mathcal{M}}_{g,r})_W$) is a sort of natural envelopping space in which to consider the various \mathcal{Q} ’s of §1. Thus, instead of asking for properness in the conventional sense, one can ask such questions as: Is some particular \mathcal{Q} closed in $(\overline{\mathcal{S}}_{g,r})_W$ (or $(\overline{\mathcal{M}}_{g,r})_W$) (Propositions 2.3, 2.4, 2.8, and 2.9; Corollary 2.6)? And, instead of considering degeneration questions in the conventional sense, one can ask: Does a particular \mathcal{Q} intersect the closure of another \mathcal{Q}' in $(\overline{\mathcal{S}}_{g,r})_W$ (Propositions 2.11 and 2.12; Corollary 2.13)? The purpose of §2 is to begin to answer such questions, and, by doing so, to shed light on the issue of compactifying moduli of Frobenius invariant indigenous bundles.

§1. Definition of the Stacks $\overline{\mathcal{Q}}_{g,r}^\Pi$

As usual, p will denote an odd prime.

§1.1. Epiperfect Schemes

Let R be a ring (commutative, with unity) of characteristic p .

Definition 1.1. We shall call R *epiperfect* if the absolute Frobenius morphism $R \rightarrow R$ is a surjection.

Let R be epiperfect. Then we may form its *perfection* as follows: We define

$$R^{\text{pf}} \stackrel{\text{def}}{=} \varprojlim R$$

where the inverse limit is that of the inverse system over \mathbf{N} that assigns to each element of \mathbf{N} a copy of R and whose transition morphisms are the Frobenius morphisms. It is trivial that R^{pf} is a *perfect ring*, i.e., that the Frobenius on R^{pf} is an isomorphism. Moreover, by projecting to the first factor in the inverse limit, we obtain a natural surjection $R^{\text{pf}} \rightarrow R$, whose kernel we denote by I^{pf} . One sees easily that R^{pf} is complete and separated with respect to the I^{pf} -adic topology. Now let us form $W(R^{\text{pf}})$, i.e., the ring of Witt vectors with coefficients in R^{pf} . Thus, we have a surjection

$$W(R^{\text{pf}}) \rightarrow R$$

obtained by composing $W(R^{\text{pf}}) \rightarrow R^{\text{pf}}$ with the surjection $R^{\text{pf}} \rightarrow R^{\text{pf}}/I^{\text{pf}} = R$. We shall denote the kernel of this surjection $W(R^{\text{pf}}) \rightarrow R$ by $I^{\text{pf}} + (p)$. Let

$$W^{\text{PD}}(R^{\text{pf}})$$

denote the result of adjoining to $W(R^{\text{pf}})$ the divided powers of the ideal $I^{\text{pf}} + (p)$. Now we make the following:

Definition 1.2. We define $B(R)$ to be the p -adic completion of $W^{\text{PD}}(R^{\text{pf}})$. We shall refer to $B(R)$ as the *universal PD-thickening* of R (cf. Lemma 1.3 below).

It is trivial that $B(R)$ is \mathbf{Z}_p -flat and p -adically complete and separated. Moreover, $B(R)$ has a natural surjection $B(R) \rightarrow R$ whose kernel has a divided power structure. In addition, $B(R)$ is equipped with a *Frobenius endomorphism*:

$$\Phi_{B(R)} : B(R) \rightarrow B(R)$$

induced by the Frobenius on $W(R^{\text{pf}})$. Finally, if R is already perfect, then $B(R) = W(R)$.

Now we have the following well-known:

Lemma 1.3. *Let C be a ring annihilated by p^n (for some $n \in \mathbf{N}$). Let $N \subseteq C$ be a PD-ideal. Let $\phi : R \rightarrow C/N$ be a morphism of rings. Then there exists a unique morphism $\phi^B : B(R) \rightarrow C$ whose composite with $C \rightarrow C/N$ is equal to the composite of $B(R) \rightarrow R$ with $\phi : R \rightarrow C/N$.*

Proof. Let $f = (f_0, f_1, \dots) \in R^{\text{pf}}$. (Thus, $f_i^p = f_{i-1}$.) Let $m \in \mathbf{N}$. Let $f'_m \in C$ be such that the image of f'_m in C/N is equal to $\phi(f_m)$. Let $f'_0 = (f'_m)^{p^m}$. Now observe that if $\alpha \in C$ and $\delta \in N$, then (since N admits divided powers) it follows that $(\alpha + \delta)^p$ is congruent to α^p modulo p . Thus,

the element $(\alpha + \delta)^{p^m} - \alpha^{p^m} \in C$ converges to zero p -adically as $m \rightarrow \infty$. In particular, if m is sufficiently large (where “sufficiently” depends only on C and N , not on f, f'_m), it follows that f'_0 is independent of the choices of m or f'_m . Thus, by mapping $[f] \mapsto f'_0 \in C$ (where $[f] \in W(R^{\text{pf}})$ is the Teichmüller representative), we obtain a set-theoretic morphism from the Teichmüller representatives of $W(R^{\text{pf}})$ to C . This set-theoretic morphism is clearly compatible with multiplication, and modulo p , it is even compatible with addition. It thus follows from a well-known argument (see, e.g., [Serre], Chapter II, §5, Proposition 9) that this set-theoretic morphism extends uniquely to a morphism of rings $W(R^{\text{pf}}) \rightarrow C$. Moreover, this morphism of rings clearly maps $I^{\text{pf}} + (p)$ into N , hence extends to a morphism $B(R) \rightarrow C$. One checks easily that this morphism has all the desired properties. \circ

Corollary 1.4. *Let R and R' be epiperfect rings. Then there is a natural one-to-one correspondence among the following three types of objects:*

- (1) *ring homomorphisms $R \rightarrow R'$;*
- (2) *ring homomorphisms $R^{\text{pf}} \rightarrow (R')^{\text{pf}}$ that map I^{pf} to $(I')^{\text{pf}}$;*
- (3) *ring homomorphisms $B(R) \rightarrow B(R')$ that map $I^{\text{pf}} + (p)$ to $(I')^{\text{pf}} + (p)$;*

Proof. Given $R \rightarrow R'$, one constructs morphisms $R^{\text{pf}} \rightarrow (R')^{\text{pf}}$ and $B(R) \rightarrow B(R')$ by functoriality. The reverse direction follows (in the case of $B(-)$) from Lemma 1.3 (where we take $C = B(R')/(p^n)$ (and N to be the kernel of $B(R')/(p^n) \rightarrow R'$) for successive n). In the case of $(-)^{\text{pf}}$ the reverse direction follows from simply passing to quotients. \circ

Let R be an epiperfect ring. For $i \geq 1$, let $R_i \stackrel{\text{def}}{=} R^{\text{pf}}/(I^{\text{pf}})^n$. Thus, R_i is also an epiperfect ring. Now the surjections $R^{\text{pf}} \rightarrow R_i$ induce a morphism:

$$\Psi : W(R^{\text{pf}}) \rightarrow \varprojlim B(R_i)$$

where the inverse limit is over the integers $i \geq 1$.

Lemma 1.5. *This morphism Ψ is an isomorphism.*

Proof. Since inverse limits commute with inverse limits, it suffices to prove that Ψ is an isomorphism when reduced modulo p^n . Moreover, by “dévissage,” it suffices to prove the following characteristic p facts:

(1) the morphism

$$R^{\text{pf}} \rightarrow \varprojlim B(R_i)_{\mathbf{F}_p}$$

(where the subscripted “ \mathbf{F}_p ” denotes reduction modulo p) is an isomorphism;

(2) the first derived functor $\mathbf{R}^1 \lim$ of the inverse system $\{B(R_i)_{\mathbf{F}_p}\}_{i \geq 1}$ is zero.

Note that there is a natural morphism $R^{\text{pf}} \rightarrow B(R)_{\mathbf{F}_p}$ (arising from the definition of $B(R)$) that factors through R_p . Let us denote by $A \subseteq B(R)_{\mathbf{F}_p}$ the subring which is the image of this morphism. Similarly, for each $B(R_i)_{\mathbf{F}_p}$, we have subrings $A_i \subseteq B(R_i)_{\mathbf{F}_p}$. Next, observe that for $j \geq 1$ sufficiently large, the elements in the kernel of $R^{\text{pf}} \rightarrow R_{i+j}$ (whose divided powers one adjoins to form $B(R_{i+j})_{\mathbf{F}_p}$) map to zero in $B(R_i)_{\mathbf{F}_p}$. In particular, it follows that for $j \geq 1$ sufficiently large, the image of $B(R_{i+j})_{\mathbf{F}_p}$ in $B(R_i)_{\mathbf{F}_p}$ is equal to the image of R^{pf} in $B(R_i)_{\mathbf{F}_p}$, hence lies inside A_i . Thus, the inverse system $\{B(R_i)_{\mathbf{F}_p}\}_{i \geq 1}$ is equivalent to the inverse system $\{A_i\}_{i \geq 1}$. But if one replaces (in (1) and (2) above) the former inverse system by the latter, then the above two facts (1) and (2) are obvious. This completes the proof. \circ

So far, we have been working with rings, but it is easy to see that all of our constructions are compatible with localization in the étale topology. Thus, we make the following definition:

Definition 1.6. If S is a scheme of characteristic p , we shall call S *epiperfect* if the Frobenius morphism $\Phi_S : S \rightarrow S$ is a closed immersion.

Note that if S is any scheme of characteristic p , then the subsheaf $\mathcal{I} \subseteq \mathcal{O}_S$ of elements whose p^{th} power is zero is a quasi-coherent sheaf of ideals. (Indeed, this follows from the fact that the Frobenius morphism $\Phi_S : S \rightarrow S$ is affine.) In particular, if S is *affine*, the cohomology of \mathcal{I} vanishes. Thus, it follows that R is epiperfect if and only if $\text{Spec}(R)$ is epiperfect. In other words, the above definition is consistent with Definition 1.1.

Moreover, if S is epiperfect, the above constructions sheafify to give sheaves of rings $\mathcal{O}_S^{\text{pf}}$, $W(\mathcal{O}_S^{\text{pf}})$, and $B(\mathcal{O}_S)$ on the étale site S_{et} of S . We shall denote by S^{pf} , $W(S^{\text{pf}})$, and $B(S)$, respectively, the corresponding ringed sites (i.e., whose structure sheaves are the preceding sheaves of rings and whose underlying sites are S_{et}).

Let S be epiperfect. Let us denote by $H_{\text{crys}}^0(S/\mathbf{Z}_p)$ the zeroth cohomology module of the structure sheaf of the *crystalline site* $\text{Crys}(S/\mathbf{Z}_p)$ of S over \mathbf{Z}_p . Note that this module has the natural structure of a ring.

Remark. Here, when we speak of the “crystalline site,” we mean the site defined (cf., e.g., [BO], the beginning of §5) using *all* PD-thickenings. Of course, there are various variants of this site obtained by imposing various conditions of nilpotence on the ideal \mathcal{I} defining the PD-thickening: e.g., the site of PD-thickenings such that there exists some n such that $\mathcal{I}^{[n]}$ (the n^{th} divided power of \mathcal{I}) is zero. We would like to remark here that one can develop the theory of this (and subsequent) chapter(s) using these modified sites instead of the standard crystalline site (as we do here), but *the end result is the same*. The reason for this is that the Frobenius invariant bundles that we are interested in all have *nilpotent p -curvature*, hence automatically define crystals both on the “standard crystalline site” (i.e., involving all PD-thickenings) as well as its various subcategories. Thus, for simplicity, we shall, for the rest of this work, stick to using the “standard crystalline site” involving all PD-thickenings.

Corollary 1.7. *Suppose that $S = \operatorname{Spec}(R)$ is affine. Then $H_{\text{crys}}^0(S/\mathbf{Z}_p) = B(R)$.*

Proof. The universal property of $B(R)$ given in Lemma 1.3 implies that we have a natural morphism $B(R) \rightarrow H_{\text{crys}}^0(S/\mathbf{Z}_p)$. On the other hand, the fact that $B(R)$ is an inverse limit of thickenings that belong to the crystalline site of S over \mathbf{Z}_p implies (by the definition of “ H_{crys}^0 ”) that we have a natural morphism $H_{\text{crys}}^0(S/\mathbf{Z}_p) \rightarrow B(R)$. Moreover, it is clear that these two morphisms are inverse to one another. This completes the proof. \circ

Note that from this point of view, the Frobenius action on $B(R)$ is precisely that induced by the Frobenius on S on the crystalline cohomology $H_{\text{crys}}^0(S/\mathbf{Z}_p)$. In particular, if $R_i \stackrel{\text{def}}{=} R^{\text{pf}}/(I^{\text{pf}})^i$, then the i^{th} power of Frobenius on $B(R_{p^i})$ factors through $B(R)$.

One familiar example of $B(R)$ is the following:

Example 1.8. Let k be a perfect field. Let $A = W(k)$. Let R be a smooth A -algebra. Let C be the p -adic completion of R . Let \overline{C} be the normalization of C in the maximal étale extension of $C \otimes_{\mathbf{Z}_p} \mathbf{Q}_p$. Then the ring $B(\overline{C}/p\overline{C})$ (in our notation) is precisely the same as the “ B_{crys} ” (i.e., ring of p -adic periods) that appears in Théorèmes 1 and 2 of §3.1 of [Font1], and almost (cf. the Remark preceding Corollary 1.7) the same as the ring $B^+(R)$ of [Falt1], §2.

§1.2. The Epiperfect Category

Let us denote by \mathfrak{Epi} the category whose objects are epiperfect schemes, and whose morphisms are scheme morphisms. Let \mathfrak{Sets} be the category of sets. Let

$$Q : \mathfrak{Epi} \rightarrow \mathfrak{Sets}$$

be a contravariant functor. Then we make the following

Definition 1.9. Let us denote by Q^{pro} the functor that assigns to $S \in \mathfrak{Obj}(\mathfrak{Epi})$ the set

$$\varprojlim Q(S)$$

where the inverse limit is taken with respect to \mathbf{N} , and the transition morphisms $Q(S) \rightarrow Q(S)$ are those induced by the Frobenius morphism on S (and the functoriality of Q). Let us denote by Q^{ind} the functor that assigns to $S \in \mathfrak{Obj}(\mathfrak{Epi})$ the set

$$\varinjlim Q(S)$$

where the direct limit is taken with respect to \mathbf{N} , and the transition morphisms $Q(S) \rightarrow Q(S)$ are those induced by the Frobenius morphism on S (and the functoriality of Q).

Thus, if $Q = \text{Spec}(A)$, where A is an \mathbf{F}_p -algebra, then Q^{pro} is represented by

$$\varinjlim A$$

where the transition morphisms in the limits are taken to be the Frobenius morphism. In particular, if A is an epiperfect ring, then Q^{pro} is represented by $\text{Spec}(A_{\text{red}})$.

Definition 1.10. If for every object $S \in \mathfrak{Obj}(\mathfrak{Epi})$, the morphism $Q(\Phi_S) : Q(S) \rightarrow Q(S)$ is an isomorphism, we shall call the functor Q *perfect*.

Note, in particular, that if Q is perfect, then $Q = Q^{\text{pro}} = Q^{\text{ind}}$.

Proposition 1.11. For any functor Q , the functors Q^{pro} and Q^{ind} are perfect.

Proof. This follows formally from the definitions. Note, moreover, when Q is represented by $\text{Spec}(A)$, the result concerning Q^{pro} follows immediately from the explicit form of Q^{pro} given above. \circ

Let $\mathfrak{P}\text{erf} \subseteq \mathfrak{E}\text{pi}$ be the full subcategory consisting of *perfect* schemes (i.e., schemes for which the Frobenius is an isomorphism). Suppose that $Q : \mathfrak{E}\text{pi} \rightarrow \mathfrak{S}\text{ets}$ is a perfect functor that converts filtered direct limits (of schemes) into inverse limits (of sets). Then it follows immediately from the definitions that if R is an epiperfect ring, then $Q(R) = Q(R^{\text{pf}})$. (Note here that since R^{pf} is an inverse limit of rings, it corresponds to a direct limit of schemes.) Thus, Q is completely determined by its restriction to $\mathfrak{P}\text{erf}$. We summarize this as follows:

Proposition 1.12. *If $Q : \mathfrak{E}\text{pi} \rightarrow \mathfrak{S}\text{ets}$ is a perfect functor such that $Q(R) = Q(R^{\text{pf}})$ (for epiperfect R), then Q is determined by its restriction to $\mathfrak{P}\text{erf}$.*

Note that the property $Q(R) = Q(R^{\text{pf}})$ can be thought of as a sort of compatibility with inverse limits.

Let Z be a smooth scheme of finite type over \mathbf{Z}_p . Let Z' be an \mathbf{F}_p -scheme equipped with a morphism $Z' \rightarrow Z$. Let us denote by Z^{pf} the functor $\mathfrak{P}\text{erf} \rightarrow \mathfrak{S}\text{ets}$ that assigns to a perfect scheme S the set of morphisms of ringed spaces $W(S) \rightarrow Z$ whose reduction modulo p is equipped with a factorization $S \rightarrow Z' \rightarrow Z$ through Z' . Then we have the following

Proposition 1.13. *The functor $Z^{\text{pf}} : \mathfrak{P}\text{erf} \rightarrow \mathfrak{S}\text{ets}$ is representable by a perfect scheme.*

Proof. First note that without loss of generality, we may assume that Z is *affine*. For $n \geq 0$, let $Z_n^{\text{pf}} : \mathfrak{P}\text{erf} \rightarrow \mathfrak{S}\text{ets}$ denote the functor that assigns to a perfect scheme S the set of morphisms $W(S) \otimes_{\mathbf{Z}_p} \mathbf{Z}/p^{n+1}\mathbf{Z} \rightarrow Z$ whose reduction modulo p is equipped with a factorization through Z' . Then Z_0^{pf} is clearly representable by the perfect scheme $(Z')^{\text{pro}}$ (where, by abuse of notation, we use the superscript “pro” over an \mathbf{F}_p -scheme X to denote the scheme given by the inverse limit

$$\lim_{\leftarrow} X$$

whose transition morphisms are the Frobenius morphisms). Now let us apply induction: That is, we shall show that Z_{n+1}^{pf} is representable by a perfect scheme under the assumption that Z_n^{pf} is. To do this, note that we have a natural morphism $\nu : Z_{n+1}^{\text{pf}} \rightarrow Z_n^{\text{pf}}$ (given by reducing modulo p^{n+1}). Moreover, since Z is smooth and affine, there exists a noncanonical section $\sigma : Z_n^{\text{pf}} \rightarrow Z_{n+1}^{\text{pf}}$ of this natural morphism. Let \mathcal{T}_n be the pull-back to Z_n^{pf} of the tangent bundle of $Z \otimes \mathbf{F}_p$ (via the natural morphism $Z_n^{\text{pf}} \rightarrow Z_0^{\text{pf}} \rightarrow Z' \rightarrow Z \otimes \mathbf{F}_p$). Then given any perfect scheme S and a morphism $\alpha : S \rightarrow Z_{n+1}^{\text{pf}}$, the difference between α and $\sigma \circ \nu \circ \alpha$

forms a section δ of $(\nu \circ \alpha)^* \mathcal{T}_n$ over S which completely determines α . Moreover, any section δ of $(\nu \circ \alpha)^* \mathcal{T}_n$ over S determines a morphism $S \rightarrow Z_{n+1}^{\text{pf}}$ which one might denote “ $\delta + (\sigma \circ \nu \circ \alpha)$.” Let $T_n \rightarrow Z_n^{\text{pf}}$ be the geometric vector bundle determined by \mathcal{T}_n . Thus, we see that Z_{n+1}^{pf} is represented by the \mathbf{F}_p -scheme T_n . Of course, T_n might not be perfect, but T_n^{pro} (which is perfect) represents the same functor on $\mathfrak{P}erf$ as T_n . This completes the induction step, and hence the proof of the Proposition. \circ

Remark. One might regard the procedure described in the proof of the above Proposition as a sort of “infinite Weil restriction of scalars” from \mathbf{Z}_p to \mathbf{F}_p .

§1.3. Epiperfect Log Schemes

Everything that we did in the preceding two subsections can be done in the logarithmic context. We begin with some definitions. Let $S^{\log} = (S, \alpha : M \rightarrow \mathcal{O}_S)$ be an integral, quasi-coherent (cf. [Kato], §2.1, 2.2, 2.3) log scheme over \mathbf{F}_p . Let us denote by $\Phi_S^{\log} : S^{\log} \rightarrow S^{\log}$ the absolute Frobenius morphism of S^{\log} .

Definition 1.14. We shall say that S^{\log} is *perfect* if Φ_S^{\log} is an isomorphism. We shall say that S^{\log} is *epiperfect* if Φ_S^{\log} is a closed immersion in the logarithmic sense (see [Kato], §3.1).

If S^{\log} is perfect (respectively, epiperfect), then it follows that S is perfect (respectively, epiperfect).

Suppose that S^{\log} is *perfect*. Then we can put a log structure on $W(S)$ as follows: Namely, we have a morphism $M \rightarrow W(\mathcal{O}_S)$ of multiplicative monoids given by mapping a section m of M to $[\alpha(m)]$ (where the brackets denote the Teichmüller representative). This morphism gives $W(S)$ a pre-log structure; taking the associated log structure ([Kato], §1.3) gives us a ringed site with log structure $W(S^{\log})$, as desired. More explicitly, the log structure of $W(S^{\log})$ is given by a morphism of monoids $W(\alpha) : W(M) \rightarrow W(\mathcal{O}_S)$, where $W(M)$ is the push-out of the following diagram:

$$\begin{array}{ccc} \mathcal{O}_S^\times & \longrightarrow & M \\ \downarrow [-1] & & \\ W(\mathcal{O}_S)^\times & & \end{array}$$

where the horizontal morphism is the natural inclusion (which exists since $\alpha : M \rightarrow \mathcal{O}_S$ defines a log structure) and the vertical morphism is given by taking the Teichmüller representative.

Now let S^{\log} be *epiperfect*. Then we may form its *perfection* as follows: We define

$$(S^{\log})^{\text{pf}} \stackrel{\text{def}}{=} \varinjlim S^{\log}$$

where the direct limit (taken in the category of log schemes) is that of the direct system over \mathbf{N} that assigns to each element of \mathbf{N} a copy of S^{\log} and whose transition morphisms are the Frobenius morphisms. It is trivial that $(S^{\log})^{\text{pf}}$ is *perfect*. By taking the inclusion via the first factor in the direct limit, we obtain a natural closed immersion $S^{\log} \hookrightarrow (S^{\log})^{\text{pf}}$. Now we can form $W((S^{\log})^{\text{pf}})$ (a ringed site with log structure) and $B(S)$ (a ringed site). Moreover, we have a natural morphism of ringed sites $B(S) \rightarrow W(S)$. Pulling back the log structure on $W((S^{\log})^{\text{pf}})$ to $B(S)$ thus gives us a log structure $B(\alpha) : B(M) \rightarrow B(\mathcal{O}_S)$ on $B(S)$. We denote the resulting logarithmic ringed site by $B(S^{\log})$.

Definition 1.15. We shall refer to $B(S^{\log})$ as the *universal log PD-thickening* of S^{\log} .

Note that we have a natural closed immersion $S^{\log} \hookrightarrow B(S^{\log})$ which makes $B(S^{\log})$ a log-PD-thickening of S^{\log} . In addition, $B(S^{\log})$ is equipped with a *Frobenius endomorphism*:

$$\Phi_{B(S^{\log})} : B(S^{\log}) \rightarrow B(S^{\log})$$

induced by the Frobenius on $W((S^{\log})^{\text{pf}})$. Finally, if S^{\log} is already perfect, then $B(S^{\log}) = W(S^{\log})$.

Lemma 1.3 generalizes immediately to the logarithmic case. (Note here that in the logarithmic version, one must replace the closed immersion of schemes $\text{Spec}(C/N) \hookrightarrow \text{Spec}(C)$ by an *exact* closed immersion of log schemes.) Corollary 1.4 and Lemma 1.5 also generalize immediately. Moreover, one has the following logarithmic version of Corollary 1.7: Let S^{\log} be epiperfect. Let us denote by $H_{\text{crys}}^0(S^{\log}/\mathbf{Z}_p)$ the zeroth cohomology module of the structure sheaf of the *logarithmic crystalline site* of S^{\log} over \mathbf{Z}_p . Note that this module has the natural structure of a p -adic ring whose formal spectrum is equipped with a pre-log structure.

Remark. The logarithmic crystalline site is discussed in [Kato], §5; however, strictly speaking in *loc. cit.*, it is assumed that the log schemes involved are *fine* (i.e., integral and coherent). Since here, our log schemes need only be integral and quasi-coherent, one must be a bit careful in dealing with the logarithmic crystalline site. In fact, however, since we will not use any deep properties of this site, it is easy to check that

everything we do here is alright, even though our log schemes are only quasi-coherent (and coherent).

Corollary 1.16. *Suppose that $S = \operatorname{Spec}(R)$ is affine. Then $H_{\text{crys}}^0(S^{\log}/\mathbf{Z}_p) = B(R)$, and the pre-log structure on $\operatorname{Spf}(B(R))$ is that arising from applying the global sections functor (on the étale site) to the étale sheaf morphism $B(\alpha) : B(M) \rightarrow B(\mathcal{O}_S)$.*

Just as in the nonlogarithmic case, if $S_0^{\log} \hookrightarrow S^{\log}$ is an exact closed immersion defined by a sheaf of ideals \mathcal{I} on S such that $\mathcal{I}^{p^i} = 0$, then the i^{th} power of the Frobenius on $B(S^{\log})$ factors through $B(S_0^{\log})$.

We denote by \mathbf{Epi}^{\log} (respectively, \mathbf{Perf}^{\log}) the category of epiperfect (respectively, perfect) log schemes. Given a functor $Q : \mathbf{Epi}^{\log} \rightarrow \mathbf{Sets}$, one can define, just as in the nonlogarithmic case, Q^{pro} and Q^{ind} , as well as the notion of a perfect functor. The functors Q^{pro} and Q^{ind} will always be perfect. Proposition 1.12 generalizes immediately to the logarithmic case. Finally, one has the following logarithmic version of Proposition 1.13:

Let Z^{\log} be a fine, smooth log scheme of finite type over \mathbf{Z}_p (where $\operatorname{Spec}(\mathbf{Z}_p)$ is endowed with the trivial log structure). Let $(Z')^{\log}$ be an \mathbf{F}_p -log scheme equipped with a morphism $(Z')^{\log} \rightarrow Z^{\log}$. Let us denote by Z^{pf} the functor $\mathbf{Perf}^{\log} \rightarrow \mathbf{Sets}$ that assigns to a perfect log scheme S^{\log} the set of morphisms of ringed log spaces $W(S^{\log}) \rightarrow Z^{\log}$ whose reduction modulo p is equipped with a factorization through $(Z')^{\log}$. Then we have the following

Proposition 1.17. *The functor $Z^{\text{pf}} : \mathbf{Perf}^{\log} \rightarrow \mathbf{Sets}$ is representable by a perfect log scheme.*

The proof of Proposition 1.17 is entirely similar to the nonlogarithmic case.

Remark. The main reason why the generalizations to the logarithmic case are so easy is because the present “epiperfect theory” is only nontrivial with respect to the “ p direction” of, say, $W(S^{\log})$ (where S^{\log} is perfect), while the log structure only plays a role in the *geometric* directions. Since these two types of directions do not interfere with each other, the logarithmic generalizations follow immediately.

§1.4. The Definition of the Stack of Quasi-Analytic Self-Isogenies

Let g and r be nonnegative integers such that $2g - 2 + r \geq 1$. Let S^{\log} be an epiperfect log scheme. Let $X^{\log} \rightarrow S^{\log}$ be an r -pointed stable

log-curve of genus g . Since $B(S^{\log})$ forms a log-PD-thickening of S^{\log} , we can consider crystals on the crystalline site $\text{Crys}(X^{\log}/B(S^{\log}))$. Note that it follows from Lemma 1.3 (and its logarithmic analogue) that crystals on $\text{Crys}(X^{\log}/B(S^{\log}))$ are equivalent to crystals on $\text{Crys}(X^{\log}/\mathbf{Z}_p)$. In what follows, we shall identify these two types of crystals by means of this equivalence. Note that Zariski locally on S , there always exists a stable r -pointed log-curve $Y^{\log} \rightarrow B(S^{\log})$ of genus g that lifts $X^{\log} \rightarrow S^{\log}$. Then a crystal on $\text{Crys}(X^{\log}/B(S^{\log}))$, valued, say, in the category of \mathbf{P}^1 -bundles, may be regarded as a \mathbf{P}^1 -bundle $P \rightarrow Y$ (over Y) equipped with a logarithmic connection ∇_P on Y^{\log} (over $B(S^{\log})$).

Let \mathcal{P} be a crystal in \mathbf{P}^1 -bundles on $\text{Crys}(X^{\log}/B(S^{\log}))$. Note that the relative Frobenius morphism $\Phi_{X/S}^{\log} : X^{\log} \rightarrow (X^{\log})^F$ on X^{\log} covers the Frobenius $\Phi_{B(S^{\log})}$ on $B(S^{\log})$, so we can pull-back the crystal \mathcal{P} via $\Phi_{X/S}^{\log}$ to obtain a new crystal $\Phi_{X/S}^*(\mathcal{P})^F$ (in \mathbf{P}^1 -bundles) on $\text{Crys}(X^{\log}/B(S^{\log}))$. Let us denote by $P_X \rightarrow X$ the \mathbf{P}^1 -bundle given by evaluating the crystal \mathcal{P} on X . Suppose that we are also given a section $\sigma : X \rightarrow P_X$, which we refer to as the *Hodge section*. Then we would like to define yet another crystal $\mathbf{F}^*(\mathcal{P})$ in \mathbf{P}^1 -bundles on $\text{Crys}(X^{\log}/B(S^{\log}))$ as follows: We shall work locally, and it will be clear in the end that everything patches together since our constructions will be canonical. First, by Zariski localizing on S , we can assume that there exists a lifting $Y^{\log} \rightarrow B(S^{\log})$ of $X^{\log} \rightarrow S^{\log}$. Thus, \mathcal{P} corresponds to a \mathbf{P}^1 -bundle with connection (P, ∇_P) on Y^{\log} (over $B(S^{\log})$). Let $U \rightarrow Y$ be a surjective étale morphism such that over U , (P, ∇_P) is the projectivization of a rank two vector bundle with connection $(\mathcal{E}, \nabla_{\mathcal{E}})$ on U^{\log} (where we equip U with the log structure pulled back from that of Y^{\log}) whose determinant is trivial. We may also assume that there exists a lifting $\Phi_U^{\log} : U^{\log} \rightarrow (U^{\log})^F \stackrel{\text{def}}{=} U^{\log} \otimes_{B(S^{\log}), \Phi_B(S^{\log})} B(S^{\log})$ of the relative Frobenius morphism in characteristic p . Let \mathcal{E}_S denote the restriction of \mathcal{E} to $U \times_{B(S)} S$. Then the given section σ corresponds to a rank one subbundle $\mathcal{L} \subseteq \mathcal{E}_S$. Note that $\Phi_U^* \mathcal{L}^F$ defines a subbundle \mathcal{M} of $(\Phi_U^* \mathcal{E}^F)_{\mathbf{F}_p}$. Indeed, this follows from the fact that:

The ideal defining the closed embedding $S \subseteq B(S)_{\mathbf{F}_p}$ is annihilated when raised to the p^{th} power.

Let $\mathcal{F} \subseteq \Phi_U^* \mathcal{E}^F$ denote the subsheaf consisting of sections whose reduction modulo p lies in \mathcal{M} . One sees easily that \mathcal{F} forms a rank two vector bundle on U and, moreover, inherits a connection $\nabla_{\mathcal{F}}$ from that of $\Phi_U^* \mathcal{E}^F$. Moreover, the projectivization of $(\mathcal{F}, \nabla_{\mathcal{F}})$ descends from U^{\log} to Y^{\log} so as to give a \mathbf{P}^1 -bundle with connection (Q, ∇_Q) on Y^{\log} . Thus, (Q, ∇_Q) defines a crystal in \mathbf{P}^1 -bundles $\mathbf{F}^*(\mathcal{P})$ on $\text{Crys}(X^{\log}/B(S^{\log}))$.

Definition 1.18. We shall refer to $\mathbf{F}^*(\mathcal{P})$ as the *renormalized Frobenius pull-back* of \mathcal{P} (with respect to the Hodge section σ). We shall refer to $\Phi_{X/S}^*(\mathcal{P})^F$ as the *naive Frobenius pull-back* of \mathcal{P} .

Now let (Π, ϖ) be a VF-pattern of period ϖ (see Chapter III, Definition 1.1). We would like to define the notion of a Π -indigenous bundle on X^{\log} as follows: We suppose that we are given a crystal in \mathbf{P}^1 -bundles

$$\mathcal{P}$$

on $\text{Crys}(X^{\log}/B(S^{\log}))$ whose restriction (P, ∇_P) to a \mathbf{P}^1 -bundle with connection on X^{\log} is indigenous. Thus, we denote the resulting Hodge section by $\sigma : X \rightarrow P$. We would like to define (inductively) a sequence

$$\mathcal{P}_0, \mathcal{P}_1, \dots$$

of crystals in \mathbf{P}^1 -bundles on $\text{Crys}(X^{\log}/B(S))$ subject to certain conditions $(*_i)$ (to be stated below). Each \mathcal{P}_i such that $\Pi(i) \neq 0$ will be equipped with a Hodge section over X . Let $\mathcal{P}_0 \stackrel{\text{def}}{=} \mathcal{P}$, and equip \mathcal{P}_0 with the Hodge section of \mathcal{P} . Now assume that \mathcal{P}_i has been defined, and satisfies the following condition:

$(*_i)$ The restriction (P_i, ∇_{P_i}) of \mathcal{P}_i to a \mathbf{P}^1 -bundle with connection on X^{\log} is crys-stable of level $\Pi(i)$.

Then (if $\Pi(i)$ is nonzero) we take the Hodge section of \mathcal{P}_i over X to be the Hodge section of the crys-stable bundle (P_i, ∇_{P_i}) (see Chapter I, Definition 3.2). Moreover,

If $\Pi(i)$ is nonzero, we define \mathcal{P}_{i+1} to be the renormalized Frobenius pull-back of \mathcal{P}_i with respect to its Hodge section. If $\Pi(i)$ is zero, then we define \mathcal{P}_{i+1} to be the naive Frobenius pull-back of \mathcal{P}_i .

Definition 1.19. If \mathcal{P} (and the bundles obtained from it as above) satisfy the conditions $(*_i)$ stated above, and, moreover, $\mathcal{P}_0 \cong \mathcal{P}_{\varpi}$, then we shall refer to \mathcal{P} as a Π -indigenous bundle of period ϖ . If \mathcal{P} is Π -indigenous of period ϖ , then we shall refer to \mathcal{P}_i as the i^{th} auxiliary bundle associated to \mathcal{P} .

Now we shall define a functor (actually a stack) which is central both to this Chapter and to this book as a whole. As above, we suppose that we are given g, r with $2g - 2 + r \geq 1$, together with a fixed VF-pattern Π of period ϖ . Our stack

$$\overline{\mathcal{Q}}_{g,r}^{\Pi} : \mathfrak{Epi}^{\log} \rightarrow \mathfrak{Spd}$$

(where \mathfrak{Spd} is the category of groupoids) is defined as follows: Given an epiperfect log scheme S^{\log} , we let $\overline{\mathcal{Q}}^{\Pi}(S^{\log})$ denote the category whose objects are pairs

$$(X^{\log} \rightarrow S^{\log}, \mathcal{P})$$

where $X^{\log} \rightarrow S^{\log}$ is an r -pointed stable log-curve of genus g and \mathcal{P} is a crystal in \mathbf{P}^1 -bundles on $\mathrm{Crys}(X^{\log}/B(S^{\log}))$ which forms a Π -indigenous bundle of period ϖ . The morphisms in the category are isomorphisms of such pairs, in the obvious sense.

Definition 1.20. We shall refer to the stack $\overline{\mathcal{Q}}_{g,r}^{\Pi}$ as the *stack of quasi-analytic self-isogenies (of r -pointed stable log-curves of genus g) of type (Π, ϖ)* .

Often, when g , r , and (Π, ϖ) are fixed throughout the discussion, we shall (for the sake of simplicity) omit them and simply write $\overline{\mathcal{Q}}^{\Pi}$, or even just \mathcal{Q} . Also, although \mathcal{Q} is, strictly speaking, a stack, it differs only very mildly from being a functor: that is, the automorphisms of a pair $(X^{\log} \rightarrow S, \mathcal{P})$ come from the automorphisms of the curve X^{\log} . Thus, by using any of the many well-known rigidifying techniques for automorphisms of a curve, one sees easily that in many cases one can, practically speaking, treat \mathcal{Q} as if it were a functor.

Remark. The reason behind the use of the term “quasi-analytic self-isogenies” is explained at the end of §1.5 of the Introduction.

§2. Deformation and Degeneration Properties of the Stacks $\overline{\mathcal{Q}}_{g,r}^{\Pi}$

In this §, we maintain the notation g , r , Π , $\mathcal{Q} = \overline{\mathcal{Q}}_{g,r}^{\Pi}$ of the preceding §.

§2.1. Lifting Properties of $\overline{\mathcal{Q}}_{g,r}^{\Pi}$

We now proceed to examine some of the basic properties of \mathcal{Q} . Note that by considering the restriction (P_i, ∇_{P_i}) of \mathcal{P}_i to X^{\log} , we obtain a collection of \mathbf{P}^1 -bundles with connection (P_i, ∇_{P_i}) on X^{\log} . By regarding (P_i, ∇_{P_i}) as the i^{th} link bundle of the “data” for an S -point of the stack $\overline{\mathcal{N}}_{g,r}^{\Pi}$ (see Chapter III, Definition 1.10; as well as the discussion as the beginning of Chapter III, §1.3, for more details on the nature on this “data”), we thus obtain a natural functor $\overline{\mathcal{Q}}_{g,r}^{\Pi}(S^{\log}) \rightarrow \overline{\mathcal{N}}_{g,r}^{\Pi}(S)$ (for every epiperfect S^{\log}). Let us endow $\overline{\mathcal{N}}_{g,r}^{\Pi}$ with the log structure obtained by pulling back the log structure of $\overline{\mathcal{M}}_{g,r}^{\log}$ via the

natural morphism $\overline{\mathcal{N}}_{g,r}^{\Pi} \rightarrow \overline{\mathcal{M}}_{g,r}$. Denote the resulting log stack by $(\overline{\mathcal{N}}_{g,r}^{\Pi})^{\log}$. Thus, we obtain a morphism of stacks

$$\overline{\mathcal{Q}}_{g,r}^{\Pi} \rightarrow (\overline{\mathcal{N}}_{g,r}^{\Pi})^{\log}$$

Let us denote (for simplicity) $(\overline{\mathcal{N}}_{g,r}^{\Pi})^{\log}$ by means of \mathcal{N} . Then the above morphism fits into a commutative diagram of natural morphisms

$$\begin{array}{ccccc} \mathcal{Q}^{\text{pro}} & \longrightarrow & \mathcal{Q} & \longrightarrow & \mathcal{Q}^{\text{ind}} \\ \downarrow & & \downarrow & & \\ \mathcal{N}^{\text{pro}} & \longrightarrow & \mathcal{N} & & \end{array}$$

Proposition 2.1. *Suppose that S^{\log} is an epiperfect log scheme such that $S = \text{Spec}(R)$; and $S_0^{\log} \subseteq S^{\log}$ is an exact closed immersion defined by a nilpotent ideal $I \subseteq R$ (i.e., there exists an n such that $I^n = 0$). Then the natural functor $\mathcal{Q}(S^{\log}) \rightarrow \mathcal{Q}(S_0^{\log})$ is fully faithful. If Π is of pure tone (see Chapter IV, Definition 2.6), then this natural functor $\mathcal{Q}(S^{\log}) \rightarrow \mathcal{Q}(S_0^{\log})$ is also essentially surjective.*

Proof. It suffices to consider the case where I is square-nilpotent, i.e., $I^2 = 0$. Let us first show that the functor is fully faithful. Let $\alpha \in \mathfrak{Ob}(\mathcal{Q}(S^{\log}))$. Let us assume that α corresponds to a curve $X^{\log} \rightarrow S^{\log}$, together with a Π -indigenous \mathcal{P} . Then we can define the crystals \mathcal{P}_i as above. Let us denote by means of a subscript “ I ” the result of base-changing objects over R by means of $R \rightarrow R/I$. Note first of all that the natural morphism $R^{\text{pf}} \rightarrow (R/I)^{\text{pf}}$ is an isomorphism, so we have a natural surjection $(R/I)^{\text{pf}} \rightarrow R$. Moreover, under this surjection, the kernel of $(R/I)^{\text{pf}} \rightarrow R/I$ is mapped to I (which is square-nilpotent, hence admits divided powers). Thus, this surjection extends to a surjection $B(R/I)_{\mathbf{F}_p} \rightarrow R$, hence a closed immersion

$$\beta : S^{\log} \hookrightarrow B(S_0^{\log})_{\mathbf{F}_p}$$

of log schemes. It thus follows that the restriction of \mathcal{P} to $\text{Crys}(X_I^{\log}/S^{\log})$ is already determined by the crystal \mathcal{P}_I on $\text{Crys}(X_I^{\log}/B(S_0^{\log}))$. Since the deformation X^{\log} of X_I^{\log} is determined by the restriction of \mathcal{P} to $\text{Crys}(X_I^{\log}/S^{\log})$ (indeed, this follows from the fact that the restriction of the crystal \mathcal{P} to X^{\log} is *indigenous*), this shows that the curve portion X^{\log} of α is determined by α_I . Now since the Frobenius on $B(S^{\log})$ factors through $B(S_0^{\log})$, the crystal \mathcal{P}_i^F on $\text{Crys}((X^{\log})^F/B(S^{\log}))$ is determined by the restriction of \mathcal{P}_i to $\text{Crys}(X_I^{\log}/B(S_0^{\log}))$. Thus, if $\Pi(i) = 0$, then the crystal $\mathcal{P}_{i+1} = \Phi_{X/S}^*(\mathcal{P}_i^F)$ on $\text{Crys}(X^{\log}/B(S^{\log}))$ is determined by the restriction of \mathcal{P}_i to

$\mathrm{Crys}(X_I^{\log}/B(S_0^{\log}))$. Moreover, if $\Pi(i) \neq 0$, then the renormalized Frobenius pull-back $\mathcal{P}_{i+1} = \mathbf{F}^*(\mathcal{P}_i)$ on $\mathrm{Crys}(X^{\log}/B(S^{\log}))$ is determined by the restriction of \mathcal{P}_i to $\mathrm{Crys}(X_I^{\log}/B(S_0^{\log}))$. Applying this to $i = \varpi - 1$, we thus obtain that \mathcal{P} is determined by α_I . This completes the proof of fully faithfulness. (Note that one justifies the fact that we acted in this proof as if \mathcal{Q} is a functor by invoking the fact that isomorphisms of objects of $\mathcal{Q}(S_0^{\log})$ lift uniquely to S^{\log} if they lift at all.)

Next let us assume that Π is of pure tone ϖ , and show essential surjectivity. Let $\alpha_I \in \mathfrak{Obj}(\mathcal{Q}(S_0^{\log}))$. Let us assume that α_I corresponds to a curve $X_I^{\log} \rightarrow S_0^{\log}$, together with a Π -indigenous \mathcal{P}_I . Just as above, by means of the closed immersion $\beta : S^{\log} \hookrightarrow B(S_0^{\log})_{\mathbf{F}_p}$, we obtain a lifting $X^{\log} \rightarrow S^{\log}$ of X_I^{\log} . Then as above, since the Frobenius on $B(S^{\log})$ admits a factorization

$$B(S^{\log}) \xrightarrow{\gamma} B(S_0^{\log}) \longrightarrow B(S^{\log})$$

(by the functoriality of $B(-)$ combined with the fact that the Frobenius on S^{\log} factors through S_0^{\log}), it follows that by base-changing by means of the first morphism γ in this factorization, we can form a crystal \mathcal{P}^F on $\mathrm{Crys}((X^{\log})^F/B(S^{\log}))$ whose restriction to $\mathrm{Crys}((X_I^{\log})^F/B(S_0^{\log}))$ is \mathcal{P}_I^F . Note also that since $\gamma_{\mathbf{F}_p}$ (i.e., the reduction of γ modulo p) factors through S^{\log} (i.e., through β), it follows from the definition of X^{\log} that \mathcal{P}^F admits a Hodge section over $X^{\log} \times_{S^{\log}, \Phi_{B(S^{\log})}} B(S^{\log})_{\mathbf{F}_p}$. Thus, by forming the renormalized Frobenius pull-back of this crystal, we obtain a crystal \mathcal{P}' on $\mathrm{Crys}(X^{\log}/B(S^{\log}))$. Now consider $(\Phi_{X/S}^*)^{\varpi-1}(\mathcal{P}')$ (which will be a crystal \mathcal{P} on $\mathrm{Crys}(X^{\log}/B(S^{\log}))$). Clearly, the restriction of \mathcal{P} to $\mathrm{Crys}(X_I^{\log}/B(S_0^{\log}))$ is equal to \mathcal{P}_I . Thus, if we restrict \mathcal{P} to $\mathrm{Crys}(X_I^{\log}/S^{\log})$ (by means of β), it follows that the resulting crystal admits a Hodge section over X^{\log} . Next, observe that if we apply the procedure just described to manufacture \mathcal{P} out of \mathcal{P}_I *twice*, we get back the same \mathcal{P} as we do when we apply this procedure *just once* (indeed, this follows from the Frobenius invariance of \mathcal{P}_I); i.e., we have

$$(\Phi_{X/S}^*)^{\varpi-1}(\mathbf{F}^*(\mathcal{P})) = \mathcal{P}$$

Thus, the data (X^{\log}, \mathcal{P}) forms an S^{\log} -valued point of \mathcal{Q} . \circ

Remark. Note that the reason that we can only prove essential surjectivity in the pure tone case is the following: Suppose, for instance, that a single period of Π contains several nonzero levels. Let i be the positive integer such that 0 and i are Π -adjacent. Then in order to obtain an S^{\log} -valued point of \mathcal{Q} , one wants the crystal $(\Phi_{X/S}^*)^{i-1}(\mathcal{P}')$ to admit a Hodge section over S . But, *a priori*, there is no reason it should do so. Moreover, attempting to change the definition of \mathcal{Q}

so as to allow different liftings X^{\log} (that agree modulo a nilpotent ideal I) at different points in the pull-back sequence does not solve the problem: Indeed, in general, the kernel of $B(R) \rightarrow R \rightarrow R/I$ does not admit divided powers, so, were one to modify the definition of \mathcal{Q} in that way, one could not even carry out the various pull-backs of crystals occurring in the definition of \mathcal{Q} .

The above Proposition implies, in particular, that the natural morphisms

$$\mathcal{Q}^{\text{pro}} \longrightarrow \mathcal{Q} \longrightarrow \mathcal{Q}^{\text{ind}}$$

induce fully faithful functors when evaluated on an epiperfect log scheme S^{\log} . Moreover, in the pure tone case, these natural morphisms are, in fact, *equivalences*.

§2.2. Representability and Affineness Properties of $\overline{\mathcal{Q}}_{g,r}^{\Pi}$

Theorem 2.2. *The stack \mathcal{Q}^{pro} is representable by a perfect algebraic log stack which is affine over $\overline{\mathcal{N}}_{g,r}^{\Pi}$. In particular, the stack \mathcal{Q}^{pro} is quasi-affine. Moreover, if Π maps every integer to χ , then \mathcal{Q}^{pro} is, in fact, affine.*

Proof. First, observe that a system of proper stable log-curves (and \mathbf{P}^1 -bundles with connections on them) lying over a filtered direct system of log schemes can always be algebrized to a proper log-curve (and \mathbf{P}^1 -bundle with connection) lying over the direct limit log scheme. It thus follows by Lemma 1.5 and Proposition 1.11 (and their logarithmic versions), plus the definition of \mathcal{Q} , that \mathcal{Q}^{pro} satisfies the hypotheses of Proposition 1.12. Thus, by Proposition 1.12, it suffices to prove that the restriction of \mathcal{Q}^{pro} , or, equivalently, of \mathcal{Q} , to $\mathfrak{P}\text{erf}^{\log}$ is representable by an algebraic log stack.

Thus, let $\alpha \in \mathfrak{Ob}(\mathcal{Q}^{\text{pro}}(S^{\log}))$, where S^{\log} is perfect. Then α corresponds to a curve $X^{\log} \rightarrow S^{\log}$, together with a crystal \mathcal{P} on $\text{Crys}(X^{\log}/B(S^{\log}))$. Note, moreover, that since the restriction of \mathcal{P} to X^{\log} is indigenous, there exists a *unique* lifting $Y^{\log} \rightarrow B(S^{\log})$ such that the evaluation (P, ∇_P) of \mathcal{P} on Y^{\log} is *indigenous* (cf. Chapter I, Proposition 4.5). Let \mathcal{S}^{\log} be the log stack (over \mathbf{Z}_p) of pairs consisting of an r -pointed stable log-curve of genus g equipped with an indigenous bundle. Note that we have a morphism of log stacks $\mathcal{N} \rightarrow \mathcal{S}^{\log}$ (where $\mathcal{N} \stackrel{\text{def}}{=} (\overline{\mathcal{N}}_{g,r}^{\Pi})^{\log}$). Thus, (P, ∇_P) defines a morphism $B(S^{\log}) = W(S^{\log}) \rightarrow \mathcal{S}^{\log}$ whose reduction modulo p factors through \mathcal{N} . Thus, we are in a position to apply Proposition 1.17. Let \mathcal{S}^{pf} be the functor of Proposition 1.17 (corresponding to the pair $(S^{\log}, \mathcal{N} \rightarrow \mathcal{S}^{\log})$).

By Proposition 1.17, we know that \mathcal{S}^{pf} is representable by a perfect algebraic log stack. In fact, it follows from the proof of Proposition 1.17 (or 1.13) that this algebraic stack is even affine over $\overline{\mathcal{N}}_{g,r}^{\Pi}$. Moreover, for every perfect S^{\log} , we have a natural fully faithful functor

$$\mathcal{Q}(S^{\log}) \rightarrow \mathcal{S}^{\text{pf}}(S^{\log})$$

On the other hand, if we take $S^{\log} = \mathcal{S}^{\text{pf}}$, then we have over $B(S^{\log}) = W(S^{\log})$, a tautological curve $Y^{\log} \rightarrow W(S^{\log})$ (pulled back from $\mathcal{S}^{\text{pf}} \rightarrow S^{\log}$), together with a tautological indigenous bundle (P, ∇_P) on Y^{\log} (also pulled back from $\mathcal{S}^{\text{pf}} \rightarrow S^{\log}$), as well as various tautological link bundles on X^{\log} (pulled back from $\mathcal{S}^{\text{pf}} \rightarrow \mathcal{N}$). We are interested in the substack of \mathcal{S}^{pf} where the following two conditions are satisfied:

- (1) the various restrictions to X^{\log} of the successive naive and renormalized Frobenius pull-backs of (P, ∇_P) coincide with the respective tautological link bundles;
- (2) the ϖ^{th} Frobenius pull-back of (P, ∇_P) coincides with (P, ∇_P) .

It is clear that this substack is represented by a *closed algebraic* (since the moduli stack of crys-stable bundles is separated algebraic – cf. Chapter I, Theorem 2.7) sub-log-stack of \mathcal{S}^{pf} . This shows that \mathcal{Q}^{pro} is representable by an algebraic log stack which is affine over $\overline{\mathcal{N}}_{g,r}^{\Pi}$, as desired. The last two sentences of the statement of the Theorem follow from Theorem 2.10 of Chapter III. This completes the proof. \bigcirc

§2.3. Embeddings of $\overline{\mathcal{Q}}_{g,r}^{\Pi}$

If X^{\log} is a log smooth algebraic log stack over \mathbf{Z}_p (endowed with the trivial log structure), let us denote by

$$X_B : \mathfrak{Epi}^{\log} \rightarrow \mathfrak{Spd}$$

the stack that assigns to an epiperfect S^{\log} the category $X^{\log}(B(S^{\log}))$, and by

$$X_W : \mathfrak{Perf}^{\log} \rightarrow \mathfrak{Spd}$$

the stack that assigns to a perfect S^{\log} the category $X^{\log}(W(S^{\log}))$. By Proposition 1.17, X_W is represented by a perfect algebraic log stack.

If $Q : \mathfrak{Epi}^{\log} \rightarrow \mathfrak{Spd}$ is a stack, we shall denote its restriction to \mathfrak{Perf}^{\log} by Q^{pf} .

As above, we fix g, r such that $2g - 2 + r \geq 1$, as well as a VF-pattern Π of period ϖ . Let \mathcal{Q} be the associated stack of quasi-analytic self-isogenies. Next, we let \mathcal{S}^{\log} be the log stack (over \mathbf{Z}_p) of r -pointed stable log-curves of genus g , equipped with an indigenous bundle. Since points of \mathcal{Q} consist of certain indigenous bundles, we obtain natural morphisms

$$\Sigma_B : \mathcal{Q} \rightarrow \mathcal{S}_B$$

and

$$\Sigma_W : \mathcal{Q}^{\text{pf}} \rightarrow \mathcal{S}_W$$

which (as follows immediately from the definition of \mathcal{Q}) are “monomorphisms” (i.e., induce fully faithful functors when evaluated on objects of \mathfrak{Epi}^{\log} or \mathfrak{Perf}^{\log}). In fact, we have the following:

Proposition 2.3. *The morphisms Σ_B and Σ_W are monomorphisms. Moreover, Σ_W is an immersion. Finally, if Π is a pre-home VF-pattern, then Σ_W is a closed immersion.*

Proof. The first statement is trivial from the definitions. Now let us show that Σ_W is an immersion. The proof is similar to that of Lemma 2.9 of Chapter III: We start by considering the tautological log-curve X^{\log} and indigenous bundle (P, ∇_P) over the perfect algebraic log stack \mathcal{S}_W . Consider the bundle $\mathbf{F}^*(P, \nabla_P)$. The reduction modulo p of this bundle is necessarily crys-stable (Proposition 1.5 of Chapter III); the condition that this bundle be of level $\Pi(1)$ is (by Theorem 3.10 of Chapter I) a condition satisfied precisely on an algebraic substack on the base stack. Continuing in this fashion, we obtain that the reductions modulo p of the various Frobenius pull-backs of (P, ∇_P) are crys-stable of the right level. The only remaining condition is that the ϖ^{th} Frobenius pull-back coincide with (P, ∇_P) ; but this condition defines a closed algebraic substack of the base. This proves that Σ_W is an immersion. When Π is a pre-home VF-pattern, the condition that the reductions modulo p of the various Frobenius pull-backs be of level $\Pi(i) = \chi$ is (by Theorem 3.10 of Chapter I) a closed condition on the base. This proves the final assertion of the Proposition. \square

If \mathcal{S}^{\log} is perfect and $n \geq 1$, let us denote by $W_n(\mathcal{S}^{\log})$ the object $W(\mathcal{S}^{\log}) \otimes \mathbf{Z}/p^{n+1}\mathbf{Z}$. For each active integer i (i.e., integer such that $\Pi(i) \neq 0$), let n_i be the unique nonnegative integer such that $i - n_i - 1$ and i are Π -adjacent (i.e., $\Pi(i - n_i - 1) \neq 0$, and, moreover, $\Pi(j) = 0$ for all integers j satisfying $i - n_i - 1 < j < i$). Let

$$\mathcal{Y} : \mathfrak{Perf}^{\log} \rightarrow \mathfrak{Spd}$$

be the stack that assigns to a perfect S^{\log} the category of objects consisting of the following data:

- (1) an r -pointed stable log-curve $X^{\log} \rightarrow S^{\log}$ of genus g ;
- (2) a crystal \mathcal{P} on $\text{Crys}(X^{\log}/W(S^{\log}))$ whose reduction modulo p is indigenous;
- (3) for each active i such that $1 \leq i \leq \varpi - 1$, a crystal in \mathbf{P}^1 -bundles \mathcal{R}_i on $\text{Crys}(X^{\log}/W_{n_i}(S^{\log}))$ whose reduction modulo p is crys-stable of level $\Pi(i)$.

One sees as in Proposition 1.17 that \mathcal{Y} is representable by a perfect algebraic log stack which, by abuse of notation, we also denote by \mathcal{Y} . Moreover, by assigning to a Π -indigenous bundle \mathcal{P} the various reductions modulo p^{n_i+1} of the \mathcal{P}_i for active i (i.e., we take these reductions for our “ \mathcal{R}_i ”), we obtain a natural morphism

$$\Delta : \mathcal{Q}^{\text{pf}} \rightarrow \mathcal{Y}$$

Now we have the following result:

Proposition 2.4. *The morphism Δ is a closed immersion.*

Proof. This proof follows a similar pattern to the that of Proposition 2.3 (with appropriate technical modifications): We start with a tautological indigenous bundle (P, ∇_P) over $Y^{\log} \rightarrow W(S^{\log})$ (for some perfect log scheme S^{\log}). If $s \in S$, then $k(s)$ is a perfect field; let us denote by s^{\log} the log scheme which is $\text{Spec}(k(s))$ equipped with the log structure induced by that of S^{\log} . We denote the reduction of Y^{\log} modulo p by $X^{\log} \rightarrow S^{\log}$. We also assume that we have been given a crystal $\mathcal{R}_i \stackrel{\text{def}}{=} (R_i, \nabla_{R_i})$ on $\text{Crys}(X^{\log}/W_{n_i}(S^{\log}))$ whose reduction modulo p is crys-stable of level $\Pi(i)$ for each active i such that $1 \leq i \leq \varpi - 1$. Suppose that $i > 0$ and 0 are Π -adjacent (so $n_i = i - 1$). Then the first condition on (P, ∇_P) that we want to consider is that $(\Phi_X^*)^{n_i}(\mathbf{F}^*(P, \nabla_P))_{\mathbf{Z}/p^{n_i+1}\mathbf{Z}}$ coincide with \mathcal{R}_i . This is a little bit more difficult than the situation of Propositions 2.3 since, although $\mathbf{F}^*(P, \nabla)_{\mathbf{F}_p}$ is always crys-stable, $(\Phi_X^*)^{n_i}(\mathbf{F}^*(P, \nabla_P))_{\mathbf{F}_p}$ need not be crys-stable in general. Thus, we are not in a situation where we obtain two S -valued points of a separated algebraic space, thus allowing us to conclude that these two points coincide precisely on some closed subscheme of S . That is to say, we need the following result:

Lemma 2.5. *Let $n \geq 1$ be an integer, and suppose that $r = 0$. Let us suppose that $\mathcal{W} = (W, \nabla_W)$ and $\mathcal{V} = (V, \nabla_V)$ are crystals on $\text{Crys}(X^{\log}/W_n(S^{\log}))$ whose reductions modulo p are crys-stable. Let $Z \subseteq S^{\text{set}}$ be the subset of the underlying set S^{set} of S consisting of points $s \in S$ such that the restrictions of $\mathcal{U} \stackrel{\text{def}}{=} (\Phi_X^*)^n \mathcal{W}$ and \mathcal{V} to $\text{Crys}(X_s^{\log}/W_n(s^{\log}))$ coincide. Then Z is closed.*

Proof. We begin with some reductions: First note that since there is only a “finite” amount of information involved, one can reduce to the case where $S^{\log} = (T^{\log})^{\text{pro}}$, where T^{\log} is a fine log scheme whose underlying scheme is of finite type over \mathbb{F}_p . Thus, we can assume that the underlying topological space of S is noetherian. Also, one sees as in Lemma 2.6 of Chapter I that the set of s where $\mathcal{U}_{\mathbb{F}_p}$ is crys-stable is a constructible subset of S . Thus, by Proposition 1.5 of Chapter I, one sees that Z is a constructible subset of S . Thus, to see that Z is closed, it suffices to see that it is closed under specialization. For this, we may reduce to the following “valuative case:” that is, the case where $S^{\log} = (T^{\log})^{\text{pro}}$; T^{\log} is a fine log scheme; and $T = \text{Spec}(B)$, where B is a discrete valuation ring. Write $S = \text{Spec}(A)$. Denote by η and s , respectively, the generic and special points of S . Let us assume that $\eta \in Z$ and show that $s \in Z$:

To do this, we would like to use the language of n -connections developed in Chapter II, §2. It is easy to check that even though there we assume that S is noetherian and S^{\log} is fine, the portion of the theory that we wish to use goes through for the present S^{\log} without any problem. Thus, ν (respectively, \mathcal{U}) defines a pre- n -connection on $\mathcal{V}_{\mathbb{F}_p}$ (respectively, $\mathcal{U}_{\mathbb{F}_p}$) (Definition 2.1 of Chapter II). Note that it is immediate from the definition of \mathcal{U} that the pre- n -connection defined by \mathcal{U} on $\mathcal{U}_{\mathbb{F}_p}$ is, in fact, an n -connection (Definition 2.2 of Chapter II). We claim that ν in fact defines an n -connection. Indeed, this follows from the fact that it defines an n -connection over η (since over η , \mathcal{U} and \mathcal{V} coincide), and the condition that a pre- n -connection be an n -connection (i.e., that the various p^i -curvatures vanish identically) is a closed condition. This proves the claim. Thus, in summary, we have an isomorphism $\text{Ad}(\mathcal{U}_{\eta})_{\mathbb{F}_p} \rightarrow \text{Ad}(\mathcal{V}_{\eta})_{\mathbb{F}_p}$ over η which is compatible with the n -connections on both bundles, and we wish to extend this isomorphism to an isomorphism over S . The rest of the proof is entirely similar to that of Proposition 1.5 of Chapter I – except that one replaces the word “connection” in *loc. cit.* by “ n -connection.”

The reason that one needs to use n -connections here is the following: Namely, when considering line bundle quotients of $\text{Ad}(\mathcal{V}_s)_{\mathbb{F}_p}$ that are stabilized just by the 0-connection, there is no problem (since $(\mathcal{V}_s)_{\mathbb{F}_p}$ is crys-stable) in concluding that such quotients have positive degree, but in the case of $\text{Ad}(\mathcal{U}_s)_{\mathbb{F}_p}$, unless one knows that a line bundle quotient is actually stabilized by the n -connection (which implies that it descends to a horizontal line bundle quotient of $\text{Ad}(\mathcal{W}_s)_{\mathbb{F}_p}$), one cannot conclude that the quotient has positive degree. \circ

Let us return to the proof of Proposition 2.4. By the above Lemma, the set of points of S where $(\Phi_X^*)^{n_i}(\mathbf{F}^*(P, \nabla_P))_{\mathbf{Z}/p^{n_i+1}\mathbf{Z}}$ coincides with \mathcal{R}_i defines a reduced closed subscheme of S . Continuing in this fashion, one sees that there is an “exact” (i.e., in the sense of log schemes) closed sub-log-scheme $Z^{\log} \subseteq S^{\log}$ over which the i^{th} Frobenius pull-back of (P, ∇_P) agrees with \mathcal{R}_i (modulo p^{n_i+1}) for active i such that $1 \leq i \leq \varpi - 1$, and the ϖ^{th} Frobenius pull-back of (P, ∇_P) is equal to (P, ∇_P) . It is clear that this Z^{\log} represents \mathcal{Q}^{pf} . This completes the proof of the Proposition. \circ

Corollary 2.6. *Let Π be a VF-pattern of pure tone (Definition 2.6 of Chapter IV). Then $\Sigma_W : \mathcal{Q}^{\text{pf}} \rightarrow \mathcal{S}_W$ is a closed immersion.*

Proof. Indeed, in this case, \mathcal{Y} coincides with \mathcal{S}_W . \circ

Remark. Recall that in the definition of $\overline{\mathcal{N}}_{g,r}^{\Pi}$ for Π whose image contains 0, we stipulated that the curve X^{\log} must be *smooth* (condition (2) of Definition 1.2 of Chapter III). At first sight, this stipulation may seem like an artificial one which means that we are only dealing with an open subset of some “more compact” complete theory. However, as one sees from Corollary 2.6, even with this smoothness condition, the resulting \mathcal{Q}^{pf} is still *complete* in the sense that it is closed in \mathcal{S}_W (which carries no conditions concerning smoothness). More generally, even \mathcal{Y} contains no condition requiring smoothness, and yet Δ is, nevertheless a *closed immersion*. The “magic” that makes this possible is Lemma 2.5, which essentially states that the specializations of Frobenius invariant indigenous bundles (as we have defined them) are again Frobenius invariant; but when the notion of Frobenius invariance involved is based on a VF-pattern that contains zeroes (hence involves *non-renormalized* Frobenius pull-backs), it follows that some of the auxiliary bundles that appear have *zero monodromy* at the nodes and marked points. On the other hand, since these auxiliary bundles are *crys-stable*, they can never (by the definition of “crys-stable” – cf. (2) and (4) of Chapter I, Definition 1.2) have zero monodromy. This contradiction is what rules out the existence of nodes and marked points in the case of VF-patterns that contain zeroes.

Suppose that S^{\log} is epiperfect, and $X^{\log} \rightarrow S^{\log}$ is an r -pointed stable log-curve of genus g . Let \mathcal{P} be a Π -indigenous bundle on X^{\log} . As usual, we form the auxiliary bundles \mathcal{P}_i . Suppose, moreover, that for each *active* i (with $0 \leq i \leq \varpi - 1$), we also give a deformation $Y_i^{\log} \rightarrow B(S^{\log})$ of X^{\log} . Then we make the following

Definition 2.7. We shall say that the $\{Y_i^{\log}\}$ *support the Π -indigenous bundle \mathcal{P}* if \mathcal{P}_i admits a Hodge section when evaluated on Y_i^{\log} .

Proposition 2.8. *Suppose that a single collection $\{Y_i^{\log}\}$ of support data supports two Π -indigenous bundles \mathcal{P} and \mathcal{P}' on X^{\log} . Then $\mathcal{P} = \mathcal{P}'$.*

Proof. If we start with \mathcal{P} , and take $\mathbf{F}^*(\mathcal{P})$, it follows easily from the definition of $\mathbf{F}^*(-)$ (cf. Lemma 2.5 of [Mzk1], Chapter III; cf. also Proposition 1.5 of Chapter III of the present work) that $\mathbf{F}^*(\mathcal{P})_{\mathbf{F}_p}$ depends only $(Y_0^{\log})_{\mathbf{Z}/p^2\mathbf{Z}}$. Suppose that i is the first active integer after 0. Then we see that $(\mathcal{P}_i)_{\mathbf{F}_p}$ is already determined by the support data, so $(\mathcal{P}_i)_{\mathbf{F}_p} = (\mathcal{P}'_i)_{\mathbf{F}_p}$. Now it follows by computing the de Rham cohomology of a crys-stable bundle of positive level (as in Proposition 1.7 of Chapter I) that the difference between two deformations (modulo a square-nilpotent ideal) of a crys-stable bundle of positive level which are supported on the same curve lies in the F^0 portion of the Hodge filtration of the first de Rham cohomology module. Applying this to the two deformations $(\mathcal{P}_i)_{\mathbf{Z}/p^2\mathbf{Z}}$ and $(\mathcal{P}'_i)_{\mathbf{Z}/p^2\mathbf{Z}}$ on $(Y_i^{\log})_{\mathbf{Z}/p^2\mathbf{Z}}$ of $(\mathcal{P}_i)_{\mathbf{F}_p}$, we thus see (just as in Lemma 2.5 of [Mzk1], Chapter III) that $\mathbf{F}^*(\mathcal{P}_i)_{\mathbf{F}_p} = \mathbf{F}^*(\mathcal{P}'_i)_{\mathbf{F}_p}$. Thus, if j is the next active integer after i , we have $(\mathcal{P}_j)_{\mathbf{F}_p} = (\mathcal{P}'_j)_{\mathbf{F}_p}$. Continuing in this way, we obtain that for all i , $(\mathcal{P}_i)_{\mathbf{F}_p} = (\mathcal{P}'_i)_{\mathbf{F}_p}$.

The same argument works modulo arbitrary powers of p , so we conclude that $\mathcal{P} = \mathcal{P}'$. For more details on what happens modulo arbitrary powers of p , we refer to the discussions of Chapter VII, §1.2, 1.3, as well as to the Pictorial Appendix (of Chapter VII). \circ

Note that conversely, if Π maps every integer to either 0 or χ , then there always exists a *unique* choice of Y_i^{\log} 's that support \mathcal{P} . Let us assume (just for the rest of this paragraph) that Π is *binary*, i.e., $\Pi(\mathbf{Z}) \subseteq \{0, \chi\}$. Let \mathcal{M} be the product (over \mathbf{Z}_p) of various copies of $(\overline{\mathcal{M}}_{g,r})_{\mathbf{Z}_p}$, indexed by the active integers i such that $0 \leq i \leq \varpi - 1$. Given a Π -indigenous bundle \mathcal{P} on a log-curve $X^{\log} \rightarrow S^{\log}$ (where S^{\log} is an object of \mathfrak{Epi}^{\log}), let $Y_i^{\log} \rightarrow B(S^{\log})$ be the unique deformation of X^{\log} over which \mathcal{P}_i becomes indigenous. Then mapping $\mathcal{P} \mapsto (Y_0^{\log}, \dots, Y_{\varpi-1}^{\log})$ induces natural morphisms

$$\Xi_B : \mathcal{Q} \rightarrow \mathcal{M}_B$$

and

$$\Xi_W : \mathcal{Q}^{\text{pf}} \rightarrow \mathcal{M}_W$$

Then, we have the following:

Proposition 2.9. *Suppose that Π is binary. Then Ξ_B is a monomorphism. If Π is the pre-home VF-pattern of period $\varpi \geq 1$, then Ξ_W is a closed immersion.*

Proof. That Ξ_B is a monomorphism is a restatement of Proposition 2.8. Let us prove that Ξ_W is a closed immersion. The argument is similar to that of Proposition 2.3. The key point is Lemma 2.5 of [Mzk1], Chapter III: Namely, if (P, ∇_P) is an indigenous bundle on $Y^{\log} \rightarrow B(S^{\log})$ (for S^{\log} epiperfect), then the reduction of $\mathbf{F}^*(P, \nabla_P)$ modulo p^j depends only on the choice of connection on the indigenous bundle modulo p^{j-2} . Thus, we start out with tautological curves $Y_i^{\log} \rightarrow W(S^{\log})$ (where $S^{\log} = \mathcal{M}_W$; $i = 0, \dots, \varpi - 1$), and choose any set of indigenous bundles $\mathcal{P}_i \stackrel{\text{def}}{=} (P_i, \nabla_{P_i})$ on Y_i^{\log} . (In fact, we may need to localize in the étale topology on \mathcal{M}_W in order to produce this set of indigenous bundles. But this is not a problem since the closed subobject (i.e., “ \mathcal{Q}^{pf} ”) of \mathcal{M}_W that we construct will be *canonical*.) Then the reductions of the $\mathbf{F}^*(\mathcal{P}_i)$ modulo p are *independent* of the choice of \mathcal{P}_i , and the condition that these reductions be indigenous is (by Theorem 3.10 of Chapter I) a *closed* condition on the base. Once we know that these reductions are indigenous, we can impose the condition that $(\mathcal{P}_{i+1})_{\mathbf{F}_p} \stackrel{\text{def}}{=} \mathbf{F}^*(\mathcal{P}_i)_{\mathbf{F}_p}$ on our arbitrary choice of indigenous bundles. Next, we consider the condition that the $\mathbf{F}^*(\mathcal{P}_i)$ ’s be indigenous modulo p^2 , which is again clearly a *closed* condition on the base. Once we know that these reductions modulo p^2 are indigenous, we can impose the condition that $(\mathcal{P}_{i+1})_{\mathbf{Z}/p^2\mathbf{Z}} \stackrel{\text{def}}{=} \mathbf{F}^*(\mathcal{P}_i)_{\mathbf{Z}/p^2\mathbf{Z}}$ on our arbitrary choice of indigenous bundles. Continuing in this fashion shows that we can realize \mathcal{Q}^{pf} as a closed algebraic substack of \mathcal{M}_W . \circ

§2.4. The Lattice of Subobjects of S_W

Let (Π, ϖ) and (Π', ϖ') be VF-patterns. Let $\mathcal{Q}' \stackrel{\text{def}}{=} \overline{\mathcal{Q}}_{g,r}^{\Pi'}$. Then by Proposition 2.3, we have immersions:

$$\Sigma_W : \mathcal{Q}^{\text{pf}} \rightarrow S_W \qquad \Sigma'_W : (\mathcal{Q}')^{\text{pf}} \rightarrow S_W$$

By abuse of notation, we can regard \mathcal{Q}^{pf} and $(\mathcal{Q}')^{\text{pf}}$ as subobjects (via these immersions) of S_W .

In this subsection, we would like to investigate how \mathcal{Q}^{pf} is positioned inside S_W relative to $(\mathcal{Q}')^{\text{pf}}$.

To do this, we introduce two orderings on the set of VF-patterns, the *left ordering* and the *right ordering*:

Definition 2.10. The *left ordering* is defined as follows: We shall say that $\Pi >_{\text{L}} \Pi'$ if there exists a negative integer n such that $\Pi(n) > \Pi'(n)$, while $\Pi(i) = \Pi'(i)$ for all negative integers $i > n$. The *right ordering* is defined as follows: We shall say that $\Pi >_{\text{R}} \Pi'$ if there exists a positive integer n such that $\Pi(n) > \Pi'(n)$, while $\Pi(i) = \Pi'(i)$ for all positive integers $i < n$.

Our first result is the following:

Proposition 2.11. *Suppose that $\Pi >_{\mathbf{L}} \Pi'$. Then inside \mathcal{S}_W , \mathcal{Q}^{pf} is disjoint from the closure of $(\mathcal{Q}')^{\text{pf}}$.*

Proof. We begin by reviewing some basic point-set topology. First, observe that if T is a topological space, $L \subseteq T$ is a locally closed subset, and $R \subseteq L$ is a subset which is *relatively closed* in L , then we have:

$$\overline{R} \cap L = R$$

(where we denote by \overline{R} the closure of R in T). Indeed, it is clear that $R \subseteq \overline{R} \cap L$, so it suffices to prove that $\overline{R} \cap L \subseteq R$. Since R is relatively closed in L , it follows that there exists a closed subset $F \subseteq T$ of T such that $R = F \cap L$. On the other hand, since $R \subseteq F$, and F is closed in T , it follows that $\overline{R} \subseteq F$, hence that $\overline{R} \cap L \subseteq F \cap L = R$, as desired. This completes the proof of the above equality.

Now let us denote by n the largest negative integer such that $\Pi(n) > \Pi'(n)$. By the point-set topology reviewed above, it suffices to find some sub-algebraic stack $\mathcal{T} \subseteq \mathcal{S}_W$ such that $\mathcal{Q}^{\text{pf}}, (\mathcal{Q}')^{\text{pf}} \subseteq \mathcal{T}$, and \mathcal{Q}^{pf} is disjoint from the closure of $(\mathcal{Q}')^{\text{pf}}$ inside \mathcal{T} . Let $Y^{\log} \rightarrow W(\mathcal{S}_W)$ be the tautological log-curve; let $X^{\log} = Y_{\mathbf{F}_p}^{\log}$. Let (P, ∇_P) be the tautological indigenous bundle on Y^{\log} . For the sake of technical simplicity, let us write that $\mathcal{P} = \mathbf{P}(\mathcal{E}, \nabla_{\mathcal{E}})$, where $(\mathcal{E}, \nabla_{\mathcal{E}})$ is a rank two vector bundle with connection on Y^{\log} whose determinant is trivial. (The reader can check easily that this technical simplification poses no essential problem.)

We would like to define, for $0 \geq i > n$, subobjects

$$\mathcal{T}_i \subseteq \mathcal{S}_W$$

as well as rank two vector bundles with connection $(\mathcal{E}_i, \nabla_{\mathcal{E}_i})$ on $Y^{\log}|_{W(\mathcal{T}_i)}$ as follows: Let $\mathcal{T}_0 \stackrel{\text{def}}{=} \mathcal{S}_W$; $(\mathcal{E}_0, \nabla_{\mathcal{E}_0}) \stackrel{\text{def}}{=} (\mathcal{E}, \nabla_{\mathcal{E}})$. Now assume that \mathcal{T}_{i+1} and $(\mathcal{E}_{i+1}, \nabla_{\mathcal{E}_{i+1}})$ have been defined, $i < 0$, and moreover:

- (1) The pre- $(-i - 1)$ -connection (defined by $(\mathcal{E}_{i+1}, \nabla_{\mathcal{E}_{i+1}})_{\mathbf{Z}/p^{-i}\mathbf{Z}}$ on the pair $(\mathcal{E}_{i+1}, \nabla_{\mathcal{E}_{i+1}})_{\mathbf{F}_p}$ is, in fact, a $(-i - 1)$ -connection (cf. Definitions 2.1 and 2.2 of Chapter II – recall that this essentially means that various “ p^j -curvatures” vanish).
- (2) There exists a crys-stable bundle $(Q_{i+1}, \nabla_{Q_{i+1}})$ on $X^{\log}|_{\mathcal{T}_{i+1}}$ and an isomorphism $(\Phi_X^*)^{-i-1} Q_{i+1} \cong \mathbf{P}(\mathcal{E}_{i+1})_{\mathbf{F}_p}$ which is “ $(-i - 1)$ -horizontal” – i.e., compatible with the

obvious $(-i-1)$ -connection on $(\Phi_X^*)^{-i-1}Q_{i+1}$ (induced by $\nabla_{Q_{i+1}}$ and the fact that this bundle is obtained by pulling back by Frobenius a total of $(-i-1)$ times) and the given $(-i-1)$ -connection (cf. (1) above) on $\mathbf{P}(\mathcal{E}_{i+1})_{\mathbf{F}_p}$.

Note that assumption (2) implies that the p^{-i} -curvature of

$(\mathcal{E}_{i+1}, \nabla_{\mathcal{E}_{i+1}})_{\mathbf{F}_p}$ is defined over all of X . (Indeed, this follows from the fact that this p^{-i} -curvature is simply the usual p -curvature of $(Q_{i+1}, \nabla_{Q_{i+1}})$.) Then we take $\mathcal{T}_i \subseteq \mathcal{T}_{i+1}$ to be the locus (i.e., reduced subscheme) where the following conditions are satisfied:

(a.) If $\Pi(i) = 0$, this p^{-i} -curvature is identically zero. In this case, we take $(\mathcal{E}_i, \nabla_{\mathcal{E}_i})$ to be $(\mathcal{E}_{i+1}, \nabla_{\mathcal{E}_{i+1}})|_{W(\mathcal{T}_i)}$. Thus, condition (1) above is satisfied when “ i ” is replaced by “ $i-1$.”

(b.) If $\Pi(i) > 0$, this p^{-i} -curvature is nilpotent with zero locus equal to the Frobenius pull-back of a $\Pi(i)$ -balanced divisor on X^{\log} . In this case, we denote by $\mathcal{L}_{i+1} \subseteq \mathcal{E}_{i+1}|_{\mathcal{T}_i}$ the unique horizontal subbundle (by which we mean that the inclusion is locally split) of rank one. Then we take $(\mathcal{E}_i, \nabla_{\mathcal{E}_i})$ to be the subsheaf of $(\mathcal{E}_{i+1}, \nabla_{\mathcal{E}_{i+1}})|_{W(\mathcal{T}_i)}$ consisting of sections whose reductions modulo p lie inside \mathcal{L}_{i+1} . Also, (as one of the defining conditions of \mathcal{T}_i) we assume that the pre- $(-i)$ -connection (defined by $(\mathcal{E}_i, \nabla_{\mathcal{E}_i})_{\mathbf{Z}/p^{-i+1}\mathbf{Z}}$ on $(\mathcal{E}_i, \nabla_{\mathcal{E}_i})_{\mathbf{F}_p}$ is, in fact, a $(-i)$ -connection. Note that in some sense, $(\mathcal{E}_i, \nabla_{\mathcal{E}_i})$ is the “inverse renormalized Frobenius pull-back” of $(\mathcal{E}_{i+1}, \nabla_{\mathcal{E}_{i+1}})$.

(c.) There exists a (Q_i, ∇_{Q_i}) as in assumption (2) above (except with respect to $(\mathcal{E}_i, \nabla_{\mathcal{E}_i})$). Note that here, *a priori*, it is necessary to take \mathcal{T}_i to be a scheme over \mathcal{T}_{i+1} (i.e., such that the morphism $\mathcal{T}_i \rightarrow \mathcal{T}_{i+1}$ is not necessarily an immersion) parametrizing crys-stable (Q_i, ∇_{Q_i}) and isomorphisms $(\Phi_X^*)^{-i}Q_i \cong \mathbf{P}(\mathcal{E}_i)$ satisfying various properties (as in (2) above). However, having done this, it is not difficult to show that this morphism $\mathcal{T}_i \rightarrow \mathcal{T}_{i+1}$ is a monomorphism, and, moreover, satisfies the valuative criterion for radimmersions (Corollary 2.13 of Chapter I), so in fact, $\mathcal{T}_i \rightarrow \mathcal{T}_{i+1}$ will be an immersion, as desired.

Thus, we obtain a sub-algebraic log stack $\mathcal{T}_i \subseteq \mathcal{T}_{i+1}$. Let $\mathcal{T} \stackrel{\text{def}}{=} \mathcal{T}_{n+1}$.

Now observe that both \mathcal{Q}^{pf} and $(\mathcal{Q}')^{\text{pf}}$ are contained in \mathcal{T} . Indeed, the restriction of $\mathbf{P}(\mathcal{E}_i, \nabla_{\mathcal{E}_i})$ to \mathcal{Q}^{pf} (respectively, $(\mathcal{Q}')^{\text{pf}}$) is none other than

$(\Phi_X^*)^{-i}(\mathcal{P}_{\varpi+i})$, where $\mathcal{P}_{\varpi+i}$ is the $(\varpi+i)^{\text{th}}$ auxiliary bundle associated to the tautological Π -indigenous (respectively, Π' -indigenous) bundle \mathcal{P} on \mathcal{Q}^{pf} (respectively, $(\mathcal{Q}')^{\text{pf}}$). Now let us consider the p^{-n} -curvature of $(\mathcal{E}_{n+1}, \nabla_{\mathcal{E}_{n+1}})_{\mathbf{F}_p}$ on $X_{\mathcal{T}}^{\log}$. Let $\mathcal{U} \subseteq \mathcal{T}$ be the locus where the zero locus of this p^{-n} -curvature forms a divisor on $X_{\mathcal{T}}^{\log}$ of degree

$$\leq 2p(\chi - \Pi(n))$$

(i.e., the degree of the Frobenius pull-back of a $\Pi(n)$ -balanced divisor). Then it is clear that $\mathcal{U} \subseteq \mathcal{T}$ is *open*, and that $\mathcal{Q}^{\text{pf}} \subseteq \mathcal{U}$, while $(\mathcal{Q}')^{\text{pf}} \cap \mathcal{U} = \emptyset$. Thus, \mathcal{Q}^{pf} is disjoint from the closure of $(\mathcal{Q}')^{\text{pf}}$ in \mathcal{T} , as desired. This completes the proof. \circ

Proposition 2.12. *Suppose that $\Pi >_{\mathbf{R}} \Pi'$. Then inside \mathcal{S}_W , $(\mathcal{Q}')^{\text{pf}}$ is disjoint from the closure of \mathcal{Q}^{pf} .*

Proof. The proof is similar to that of Proposition 2.11, only somewhat technically *easier*. We remark that the reason that the present case is easier than that of Proposition 2.11 is that whereas in the proof of Proposition 2.11, we had to “*peel off*” successive renormalized and naive Frobenius pull-backs, in the present case, we need merely *apply* successive renormalized and naive Frobenius pull-backs.

As above, we let $Y^{\log} \rightarrow W(\mathcal{S}_W)$ be the tautological log-curve; let $X^{\log} = Y_{\mathbf{F}_p}^{\log}$. Let \mathcal{P} be the tautological indigenous bundle on Y^{\log} . Let n be the smallest positive integer such that $\Pi(n) > \Pi'(n)$. Then let us define, for nonnegative $i < n$, subobjects

$$\mathcal{T}_i \subseteq \mathcal{S}_W$$

together with \mathcal{P}_i on $Y_{W(\mathcal{T}_i)}^{\log}$ as follows: Let $\mathcal{T}_0 \stackrel{\text{def}}{=} \mathcal{S}_W$; $\mathcal{P}_0 \stackrel{\text{def}}{=} \mathcal{P}$. Assume that \mathcal{T}_{i-1} and \mathcal{P}_{i-1} have been defined, and that the reduction modulo p of \mathcal{P}_{i-1} on $X_{\mathcal{T}_{i-1}}^{\log}$ is *crys-stable* of level $\Pi(i-1)$. Let

$$\mathcal{T}_i \subseteq \mathcal{T}_{i-1}$$

be the reduced subobject where:

If $\Pi(i-1) = 0$ (respectively, $\Pi(i-1) > 0$), then $\Phi_X^*(\mathcal{P}_{i-1})_{\mathbf{F}_p}$ (respectively, $\mathbf{F}^*(\mathcal{P}_{i-1})_{\mathbf{F}_p}$) is *crys-stable* of level $\Pi(i)$.

Let \mathcal{P}_i be the restriction to $Y_{W(\mathcal{T}_i)}^{\log}$ of the respective (renormalized or naive) Frobenius pull-back of \mathcal{P}_{i-1} . Thus, $\mathcal{T}_i \subseteq \mathcal{T}_{i-1}$ is a sub-algebraic log stack. Let $\mathcal{T} \stackrel{\text{def}}{=} \mathcal{T}_{n-1}$. Note that $\mathcal{Q}^{\text{pf}}, (\mathcal{Q}')^{\text{pf}} \subseteq \mathcal{T}$.

Let $\mathcal{F} \subseteq \mathcal{T}$ be the locus where

- (1) If $\Pi(n-1) = 0$, then $\Phi_X^*(\mathcal{P}_{n-1})_{\mathbf{F}_p}$ is crys-stable (this much is automatic – since $(\mathcal{P}_{n-1})_{\mathbf{F}_p}$ is crys-stable of level $0 = \Pi(n-1)$), and defines a point (in the moduli space of crys-stable bundles) which is in the closure of the locus of crys-stable bundles of level $\Pi(n)$.
- (2) If $\Pi(n-1) > 0$, $\mathbf{F}^*(\mathcal{P}_{n-1})_{\mathbf{F}_p}$ is crys-stable (this much is automatic – cf. Proposition 1.5 of Chapter III), and defines a point (in the moduli space of crys-stable bundles) which is in the closure of the locus of crys-stable bundles of level $\Pi(n)$.

Then clearly $\mathcal{F} \subseteq \mathcal{T}$ is closed, and $\mathcal{Q}^{\text{pf}} \subseteq \mathcal{F}$. Moreover, by Theorem 3.10 of Chapter I, it follows that $(\mathcal{Q}')^{\text{pf}}$ does not intersect \mathcal{F} . This completes the proof. \circ

Corollary 2.13. *Let $\mathcal{H} \subseteq \mathcal{S}_W$ be a union of any finite set of $(\mathcal{Q}^\Pi)^{\text{pf}}$'s where Π is pre-home. Let $\mathcal{K} \subseteq \mathcal{S}_W$ be a union of any finite set of $(\mathcal{Q}^\Pi)^{\text{pf}}$'s where Π is not pre-home. Then $\mathcal{H} \subseteq \mathcal{S}_W$ is closed and disjoint from the closure of \mathcal{K} .*

Proof. If Π is pre-home and Π' is not pre-home, then $\Pi >_{\text{L}} \Pi'$. Now apply Propositions 2.3 and 2.11. \circ

Remark. Another way to put the above Corollary is that the union of pre-home patterns is (in some sense) closed and isolated from the other \mathcal{Q}^{pf} 's in \mathcal{S}_W .

Chapter VII: The Generalized Ordinary Theory

§0. Introduction

In Chapter VI, we introduced the stacks \mathcal{Q} as universal stacks of Frobenius invariant indigenous bundles. It would be nice if we could make these stacks explicit in general, but at the present time this seems to be too technically difficult. Thus, just as in [Mzk1], we end up having to restrict to a certain “ordinary” (i.e., more precisely, Π -ordinary, as defined below) locus. The condition of being Π -ordinary is manifestly open. Thus, over this open substack of \mathcal{Q} , we obtain (Theorem 1.8) a rather explicit theory involving Frobenius liftings which is reminiscent of, but substantially generalizes the theory of [Mzk1] (cf. especially Theorem 2.8 of [Mzk1], Chapter III). Unfortunately, this pleasant explicitness with respect to the *lifting* side of things is gained at the price of restricting ourselves to an *open* substack of \mathcal{Q} . That is to say, one can ask what happened to our goal of also trying to obtain *compact* moduli stacks of Frobenius invariant indigenous bundles. Indeed, if one does not obtain “compact” (in some appropriate sense) moduli stacks, there always remains the suspicion that one’s theory is not *complete*, i.e., one can generalize the theory a bit beyond what one already has by looking at *degenerations* of the objects whose moduli one’s theory already describes.

One answer to this dilemma is discussed in §2 (Theorem 2.11): Namely, we show that in fact, at least under a certain restriction on the VF-pattern Π , the Π -ordinary locus is not only open, *but also* “ ω -closed” (cf. Definitions 2.5 and 2.10) inside \mathcal{Q} . Roughly speaking, “ ω -closed” means that at least “as far as the differentials are concerned,” it is as if the Π -ordinary locus is closed. In fact, among one-dimensional schemes or schemes which are perfectionizations of schemes which are locally of finite type, ω -closedness already implies closedness (cf. Propositions 2.7, 2.8). Thus:

At least if one is satisfied with perfections of locally finite type subobjects of \mathcal{Q} or $(\overline{\mathcal{S}}_{g,r})_W$ (cf. Chapter VI) as one's "enveloping universe," *the Π -ordinary theory (for binary Π) is already a compact or "complete" theory.*

We remark that in an earlier version of this work ([Mzk5,6]), we announced that we had in fact shown that the Π -ordinary locus is not just ω -closed, but closed (in the usual sense). Closer inspection of the proof, however, revealed a gap, which, in essence, amounts to the statement that the proof the author envisioned only implies ω -closedness, not closedness. The difference between ω -closedness and closedness, is that ω -closedness is, essentially, closedness within all possible compactifications (which are the "perfections" of an object) whose canonical sheaf ω is *bounded* – i.e., such that multiplication by a non-invertible function on ω induces a non-invertible endomorphism of ω . A more geometric way to think of the condition that ω be bounded is that it means that the object in question admits a notion of *volume* which is locally *finite*. Thus, any sort of compactification that conforms to one's typical intuition of "variety" (e.g., any scheme which is locally of finite type over the base field) will have a bounded ω .

Finally, §3 is concerned with the important question of existence. That is to say, no matter how elegant a theory one obtains concerning the Π -ordinary locus of \mathcal{Q} , this theory is meaningless if it turns out that the Π -ordinary locus is, in fact, empty! In [Mzk1] (cf. Corollary 3.8 of [Mzk1], Chapter II), it was relatively trivial to show that the ordinary locus is nonempty. In the present generalized context, however, things are not so trivial. However, in §3, by building on the extensive results of Chapters II and IV, we show that in a wide variety of cases (Theorems 3.1 and 3.7), *the Π -ordinary locus is nonempty*. In addition, in the spiked case, we show that *the "very Π -ordinary locus" is nonempty* (Theorem 3.7). The "very Π -ordinary locus" (an open substack of the Π -ordinary locus) is introduced because on this locus it is possible to analyze the canonical Frobenius liftings of Theorem 1.8 in greater detail (Theorem 3.8). Moreover, as noted in Chapters II and IV, the machinery developed there allows the reader to construct further examples of the Π -ordinary theory by reducing the issue of constructing such examples to essentially *combinatorial* (as opposed to *arithmetic algebro-geometric*) problems.

§1. The Π -Ordinary Locus

§1.1. The Frobenius Action on the Crystalline Cohomology

Let S^{\log} be perfect; let $f^{\log} : X^{\log} \rightarrow S^{\log}$ be an r -pointed stable log-curve of genus g . Let Π be a VF-pattern of period ϖ . Let \mathcal{P}

be a crystal on $\text{Crys}(X^{\log}/W(S^{\log}))$ which forms a Π -indigenous bundle of period ϖ (Definition 1.19 of Chapter VI). Let $\mathcal{P}_1, \dots, \mathcal{P}_\varpi = \mathcal{P}$ be the associated auxiliary bundles. Note that for $i, n \geq 0$, we may form the relative crystalline cohomology modules $\mathbf{R}^n f_{\text{cr},*}(\text{Ad}(\mathcal{P}_i))$ (which are vector bundles on $W(S)$). By Proposition 1.6 of Chapter I, it follows that the $\mathbf{R}^n f_{\text{cr},*}(\text{Ad}(\mathcal{P}_i))$ are zero unless $n = 1$. Thus, we shall use the notation

$$\mathcal{H}(\mathcal{P}_i) \stackrel{\text{def}}{=} \mathbf{R}^1 f_{\text{cr},*}(\text{Ad}(\mathcal{P}_i))$$

in the following discussion. By Proposition 1.6 of Chapter I, each $\mathcal{H}(\mathcal{P}_i)$ is a vector bundle of rank $2(3g-3+r)$ on $W(S)$ which is naturally *self-dual* (since in the present context, all the radii are zero). Next, observe that the reduction $\mathcal{H}(\mathcal{P}_i)_{\mathbf{F}_p}$ of each cohomology module modulo p is equipped with a Hodge filtration $F^j(\mathcal{H}(\mathcal{P}_i)_{\mathbf{F}_p})$ whose only possibly nonzero subquotients lie in degrees $-1, 0, 1$, and 2 . Moreover, each of these subquotients forms a *vector bundle* on S . Indeed, for F^{-1}/F^0 and F^2 (and hence also for F^0/F^2), this follows from Proposition 1.7 of Chapter I; for F^0/F^1 , this follows from Lemma 3.8 of Chapter I; and for F^1/F^2 , this follows from the fact that F^1/F^2 is the kernel of the surjection of vector bundles $F^0/F^2 \rightarrow F^0/F^1$. Let us denote by

$$\mathcal{H}'(\mathcal{P}_i) \subseteq \mathcal{H}(\mathcal{P}_i)$$

the subsheaf of $\mathcal{H}(\mathcal{P}_i)$ consisting of sections whose reduction modulo p lies in $F^0(-)$. Then, as usual, one obtains a natural Frobenius action on the crystalline cohomology

$$\mathcal{F}_i : \Phi_{W(S)}^* \mathcal{H}'(\mathcal{P}_i) \rightarrow \mathcal{H}(\mathcal{P}_{i+1})$$

We would like to consider the reduction of this Frobenius action modulo p . If we reduce \mathcal{F}_i modulo p , we obtain a morphism

$$(\mathcal{F}_i)_{\mathbf{F}_p} : \Phi_S^* \mathcal{H}'(\mathcal{P}_i)_{\mathbf{F}_p} \rightarrow \mathcal{H}(\mathcal{P}_{i+1})_{\mathbf{F}_p}$$

Now observe that the $F^1(-)$ -portion of the Hodge filtration on $\mathcal{H}(\mathcal{P}_i)_{\mathbf{F}_p}$ induces a submodule $F^1(-)$ on $\mathcal{H}'(\mathcal{P}_i)_{\mathbf{F}_p}$. Indeed, at first, it may look like one needs to choose a lifting of $X^{\log} \rightarrow S^{\log}$ to a curve over $W(S^{\log})_{\mathbf{Z}/p^2\mathbf{Z}}$, which would define a lifting of $F^1(\mathcal{H}(\mathcal{P}_i)_{\mathbf{F}_p})$ to a $F^1(\mathcal{H}(\mathcal{P}_i)_{\mathbf{Z}/p^2\mathbf{Z}})$, in order to consider the “image” of $F^1(-)$ in $\mathcal{H}'(\mathcal{P}_i)_{\mathbf{F}_p}$. However, one sees immediately that by “Griffiths transversality” (i.e., different liftings only deform $F^1(-)$ into $p \cdot F^0(-)$, whose image vanishes in $\mathcal{H}'(\mathcal{P}_i)_{\mathbf{F}_p}$) all liftings give rise to the same image in $\mathcal{H}'(\mathcal{P}_i)_{\mathbf{F}_p}$. Thus, in summary, it makes sense to speak of $F^1(-)$ of $\mathcal{H}'(\mathcal{P}_i)_{\mathbf{F}_p}$ or $\mathcal{H}(\mathcal{P}_{i+1})_{\mathbf{F}_p}$;

moreover, these $F^1(-)$'s are both vector bundles of rank $3g - 3 + r$ (cf. Lemma 3.8 of Chapter I).

Let us denote by $\Theta(\mathcal{P}_i)$ (respectively, $\Theta'(\mathcal{P}_i)$) the quotient of $\mathcal{H}(\mathcal{P}_i)_{\mathbf{F}_p}$ (respectively, $\mathcal{H}'(\mathcal{P}_i)_{\mathbf{F}_p}$) by $F^1(-)$. Thus, both $\Theta(\mathcal{P}_i)$ and $\Theta'(\mathcal{P}_i)$ are vector bundles on S of rank $3g - 3 + r$. Moreover, $\Theta(\mathcal{P}_i)$ gets an induced Hodge filtration from that of $\mathcal{H}(\mathcal{P}_i)_{\mathbf{F}_p}$, so we have an exact sequence

$$0 \rightarrow (F^0/F^1)(\mathcal{H}(\mathcal{P}_i)_{\mathbf{F}_p}) \rightarrow \Theta(\mathcal{P}_i) \rightarrow (F^{-1}/F^0)(\mathcal{H}(\mathcal{P}_i)_{\mathbf{F}_p}) \rightarrow 0$$

while $\Theta'(\mathcal{P}_i)$ gets an “inverted” (or “conjugate”) Hodge filtration:

$$0 \rightarrow (F^{-1}/F^0)(\mathcal{H}(\mathcal{P}_i)_{\mathbf{F}_p}) \rightarrow \Theta'(\mathcal{P}_i) \rightarrow (F^0/F^1)(\mathcal{H}(\mathcal{P}_i)_{\mathbf{F}_p}) \rightarrow 0$$

Note by the usual yoga of Hodge filtrations and Frobenius actions on crystalline cohomology (i.e., that when Frobenius acts on crystalline cohomology, it sends $F^1(-)$ to 0 modulo p), it follows that the morphism $(\mathcal{F}_i)_{\mathbf{F}_p}$ factors through $\Theta'(\mathcal{P}_i)$, so we obtain a morphism

$$\Phi_i : \Phi_S^* \Theta'(\mathcal{P}_i) \rightarrow \Theta(\mathcal{P}_{i+1})$$

from $(\mathcal{F}_i)_{\mathbf{F}_p}$.

Definition 1.1. We shall say that \mathcal{P} is Π -ordinary if Φ_i is an isomorphism for all $i \in \mathbf{Z}$.

Remark 1. More generally, if we only assume that S^{\log} is epiperfect, and \mathcal{P} is a crystal on $\text{Crys}(X^{\log}/B(S^{\log}))$ which forms a Π -indigenous bundle of period ϖ , then we can extend the above definition as follows: Let T^{\log} be the perfect scheme obtained by taking the inverse limit $(S^{\log})^{\text{pro}}$ (i.e., \mathcal{O}_T is the quotient of \mathcal{O}_S by its subsheaf of nilpotent elements). Let $X_T^{\log} \rightarrow T^{\log}$ (respectively, \mathcal{P}_T) be the restriction of X^{\log} (respectively, \mathcal{P}) to T^{\log} (respectively, $\text{Crys}(X_T^{\log}/W(T^{\log}))$). Then we shall call \mathcal{P} Π -ordinary if \mathcal{P}_T is Π -ordinary in the sense of Definition 1.1.

Remark 2. Note that when (Π, ϖ) is the home VF-pattern, the notion of ordinarieness here is compatible with that introduced in [Mzk1], Chapter II, Definition 3.1.

Clearly, there is a stack \mathcal{Q}^{ord} of Π -ordinary Π -indigenous bundles, which forms an open substack of \mathcal{Q} (the stack of Π -indigenous bundles of period ϖ). Similarly, one has an open substack $\mathcal{Q}^{\text{pro,ord}} \subseteq \mathcal{Q}^{\text{pro}}$. Since \mathcal{Q}^{pro} is representable by a perfect algebraic log stack (Theorem 2.2 of Chapter VI), we have the following:

Proposition 1.2. *For a VF-pattern Π of period ϖ , the stack $\mathcal{Q}^{\text{pro,ord}}$ of Π -ordinary Π -indigenous bundles of period ϖ is representable by a perfect algebraic log stack, which is an open algebraic substack of \mathcal{Q}^{pro} .*

§1.2. Interpretation of the Condition of Π -Ordinariness

We continue with the notation of the preceding subsection. Suppose that the restriction of the crystal \mathcal{P}_i to $\text{Crys}(X^{\log}/S^{\log})$ is a crystal bundle (P_i, ∇_{P_i}) of level $\Pi(i)$ on X^{\log} . Let

$$H_i = \mathbf{R}^1 f_{\text{DR},*} \text{Ad}(P_i)$$

Thus, H_i is a vector bundle of rank $2(3g-3+r)$ on S . In the notation of the preceding subsection, $H_i = \mathcal{H}(\mathcal{P}_i)_{\mathbf{F}_p}$. We would first like to review the “meaning” of the various subquotients of the Hodge filtration on H_i .

In the case $\Pi(i) = 0$, one sees easily that we have

$$F^{-1} = F^0; \quad F^2 = 0; \quad F^0/F^1 \cong (F^1)^\vee$$

and, moreover,

$$F^0/F^1 = \mathbf{R}^1 f_* \text{Ad}(P_i)$$

Thus, if one thinks of H_i as representing infinitesimal deformations of (P_i, ∇_{P_i}) , the projection $H_i \rightarrow F^0/F^1$ corresponds to “just looking at the \mathbf{P}^1 -bundle deformation,” i.e., forgetting the connection, while the injection $F^1 \subseteq H_i$ corresponds to deformations obtained by fixing the \mathbf{P}^1 -bundle $P_i \rightarrow X$ and deforming the connection ∇_{P_i} .

On the other hand, in the case $\Pi(i) > 0$, let $D_i \subseteq X$ be the Kodaira-Spencer locus of (P_i, ∇_{P_i}) . Let $\mathcal{L} = \omega_{X/S}^{\log}(-D_i)$. By Proposition 1.7 of Chapter I, the filtration on H_i has the following property: $F^{-1}/F^0 = \mathbf{R}^1 f_* \mathcal{L}^{-1}$, and moreover, the projection $H_i \rightarrow F^{-1}/F^0$ corresponds to looking at the obstruction to lifting the Hodge section. Thus, $F^0 \subseteq H_i$ corresponds to deformations for which the Hodge section lifts. Moreover, the projection $F^0 \rightarrow F^0/F^1$ corresponds precisely to looking at the induced deformation of the Kodaira-Spencer locus D_i (cf. Lemma 3.8 of Chapter I).

Let $\Theta'_i \stackrel{\text{def}}{=} \Theta'(\mathcal{P}_i)$; let $\Theta_i^F \stackrel{\text{def}}{=} \Phi_S^* \Theta'_i$. Recall that \mathcal{F}_i induces a morphism

$$\Xi_i : \Theta_i^F \rightarrow H_{i+1}$$

We would like to interpret this morphism. If $\Pi(i) = 0$, then $\Theta_i^F = \Phi_S^*(F^0/F^1(H_i))$. Thus, relative to the formation of the crys-stable bundle $\Phi_X^*(P_i)$, the morphism Ξ_i corresponds in the obvious way to associating to the pull-back of the deformation of P_i the corresponding deformation of $\Phi_X^*(P_i, \nabla_{P_i})$. Now, assume that $\Pi(i) > 0$. Then one has an exact sequence

$$0 \rightarrow \Phi_S^*(F^{-1}/F^0) \rightarrow \Theta_i^F \rightarrow \Phi_S^*(F^0/F^1) \rightarrow 0$$

Relative to the formation of $\mathbf{F}^*(\mathcal{P}_i)_{\mathbf{F}_p}$, this sequence may be interpreted as follows: The crys-stable bundle $\mathbf{F}^*(\mathcal{P}_i)_{\mathbf{F}_p}$ will be a bundle of the sort discussed in Proposition 1.5 of Chapter III. The projection to $\Phi_S^*(F^0/F^1)$ in the above exact sequence then corresponds to the deformation of the p -curvature locus of $\mathbf{F}^*(\mathcal{P}_i)_{\mathbf{F}_p}$. On the other hand, the inclusion of $\Phi_S^*(F^{-1}/F^0)$ corresponds to the situation where the p -curvature locus remains fixed, but one deforms the section of the torsor “ \mathcal{A} ” of Proposition 1.5 of Chapter III. (Indeed, note that in the notation of Proposition 1.5 of Chapter III, the bundle $\Phi_S^*(F^{-1}/F^0)$ is simply “ $\Theta_{D_i^F}$ ” – i.e., what was called “ Θ_D ” in *loc. cit.* applied to the case where “ D ” is taken to be the Frobenius transform $D_i^F \subseteq X^F$ of the divisor $D_i \subseteq X$.)

Write \mathcal{N} (respectively, \mathcal{N}^s) for $(\overline{\mathcal{N}}_{g,r}^\Pi)^{\log}$ (respectively, $(\overline{\mathcal{N}}_{g,r}^{\Pi,s})^{\log}$) (where the log structure is that obtained by pulling back the log structure on $\overline{\mathcal{M}}_{g,r}^{\log}$) – i.e., the *shifted* VF-stack of Definition 1.11 of Chapter III). Observe that \mathcal{N} and \mathcal{N}^s represent the same stack on $\mathfrak{P}\mathrm{erf}^{\log}$.

Now let $\alpha \in \mathcal{N}(S^{\log})$ be the S^{\log} -valued point of \mathcal{N} defined by the reductions modulo p of $\mathcal{P}_0, \dots, \mathcal{P}_{\varpi-1}$. Now I *claim* the following:

(*) *The Π -indigenous bundle \mathcal{P} is Π -ordinary if and only if the natural morphism $\overline{\mathcal{N}}_{g,r}^{\Pi,s} \rightarrow \overline{\mathcal{M}}_{g,r}$ is étale at α .*

Indeed, first let us observe that the Π -ordinariness of \mathcal{P} implies that $\overline{\mathcal{N}}_{g,r}^{\Pi,s} \rightarrow \overline{\mathcal{M}}_{g,r}$ is étale at α . To see this, we must show that given a deformation of the underlying curve X^{\log} , there exists a unique deformation of all the data (cf. the discussion at the end of §1.3) that make up a point of $\overline{\mathcal{N}}_{g,r}^{\Pi,s}$. First of all, note that since the various curves X_i^{\log} (of *loc. cit.*) are isomorphic either to X^{\log} or to various Frobenius conjugates of X^{\log} , it follows that the X_i^{\log} all deform uniquely. Thus, it remains to see that all the pairs (P_i, ∇_{P_i}) deform uniquely. We do this by descending induction on i , starting with $i = 0$. Suppose that the deformation of the Kodaira-Spencer locus $D_i \subseteq X_i$ (respectively, \mathbf{P}^1 -bundle P_i) – if $\Pi(i) > 0$ (respectively, if $\Pi(i) = 0$) – has already been determined uniquely. (Note that this is the case for $i = 0$.) Then note that the fact that Φ_{i-1} is an isomorphism, plus the fact that the deformation $(\in H_i)$ that we are looking for

is necessarily in the image of Ξ_{i-1} , *imply* (by the above discussion) that a pair (P_i, ∇_{P_i}) with the desired properties exists and is uniquely determined by D_i (respectively, P_i) if $\Pi(i) > 0$ (respectively, if $\Pi(i) = 0$). Moreover, if $\Pi(i-1) > 0$ (respectively, $\Pi(i-1) = 0$), then taking the p -curvature locus of (P_i, ∇_{P_i}) (respectively, the unique \mathbf{P}^1 -bundle on X_{i-1} whose sections are the horizontal sections of P_i) shows that D_{i-1} (respectively, P_{i-1}) is uniquely determined, thus allowing us to continue the induction. This completes the verification of the “only if” part of the claim.

Next, let us verify the “if” part of the claim. Thus, we assume that $\overline{\mathcal{N}}_{g,r}^{\Pi,s} \rightarrow \overline{\mathcal{M}}_{g,r}$ is étale at α . We would like to show that all the Φ_i are isomorphisms. Note that it suffices to prove this when S is the spectrum of a field, so in the rest of this paragraph, we shall assume this (for the sake of simplicity). Now let i be the smallest nonnegative integer for which Φ_i is *not* an isomorphism. Since we are working over a field, this means that Φ_i is not injective. Thus, we may construct a deformation of α as follows: We take the trivial deformation of all the X_j^{\log} (for $j \in \mathbf{Z}$). Also, we take the trivial deformation of (P_j, ∇_{P_j}) for all integers j such that $i+1 < j \leq \varpi$. Next, we deform $(P_{i+1}, \nabla_{P_{i+1}})$ by a nonzero element of H_{i+1} which is in the image (under Ξ_i) of the kernel of Φ_i . Now by taking – if $\Pi(i) > 0$ (respectively, $\Pi(i) = 0$) – the p -curvature locus of $(P_{i+1}, \nabla_{P_{i+1}})$ (respectively, the unique \mathbf{P}^1 -bundle on X_i whose sections are the horizontal sections of P_{i+1}), we obtain a $D_i \subseteq X_i$ (respectively, P_i on X_i). Now the argument of the preceding paragraph (plus the fact that Φ_j is an isomorphism if $0 \leq j < i$) allows us to construct (P_j, ∇_{P_j}) for j in the range $0 < j \leq i$ that satisfy the necessary conditions. This gives us a deformation of α which induces the trivial deformation X^{\log} , thus contradicting the assumption that $\overline{\mathcal{N}}_{g,r}^{\Pi,s} \rightarrow \overline{\mathcal{M}}_{g,r}$ is étale at α . This completes the proof of the claim (*).

We summarize the above discussion as follows:

Definition 1.3. We shall denote by $(\overline{\mathcal{N}}_{g,r}^{\Pi,s})^{\log, \text{ord}}$ the open log substack where $\overline{\mathcal{N}}_{g,r}^{\Pi,s} \rightarrow \overline{\mathcal{M}}_{g,r}$ is étale. We shall refer to $(\overline{\mathcal{N}}_{g,r}^{\Pi,s})^{\text{ord}}$ as the *ordinary locus* of $\overline{\mathcal{N}}_{g,r}^{\Pi,s}$.

Lemma 1.4. Under the natural morphism $\mathcal{Q} \rightarrow \mathcal{N}^s$ (cf. the diagram preceding Proposition 2.1 of Chapter VI), the open substack $\mathcal{Q}^{\text{ord}} \subseteq \mathcal{Q}$ is the inverse image of $\mathcal{N}^{s, \text{ord}} \subseteq \mathcal{N}^s$.

Proof. This is a restatement of the claim (*) discussed above. \square

Lemma 1.5. Suppose that we are given a morphism $S^{\log} \rightarrow \mathcal{N}^{s, \text{ord}}$. Then there exists a unique morphism $S^{\log} \rightarrow \mathcal{Q}^{\text{pro}}$ whose composite with $\mathcal{Q}^{\text{pro}} \rightarrow \mathcal{N}^s$ is the morphism obtained by composing the given morphism $S^{\log} \rightarrow \mathcal{N}^{s, \text{ord}}$ with the natural inclusion $\mathcal{N}^{s, \text{ord}} \hookrightarrow \mathcal{N}^s$.

Proof. We wish to construct a Π -indigenous bundle \mathcal{P} on $\text{Crys}(X^{\log}/W(S^{\log}))$ (as well as auxiliary bundles $\mathcal{P}_1, \dots, \mathcal{P}_{\varpi-1}$). Moreover, since one knows what the projection $S^{\log} \rightarrow \mathcal{N}^s$ should be, it follows that the $(\mathcal{P}_i)_{\mathbf{F}_p}$ are all determined. Thus, one is reduced to a question of the existence (and uniqueness) of appropriate liftings. We begin by introducing some useful terminology:

- (1) Let \mathcal{A} and \mathcal{B} be crystals in \mathbf{P}^1 -bundles on $\text{Crys}(X^{\log}/W(S^{\log}))$ whose reductions modulo p are crystal-stable bundles of some level. Then we shall say that \mathcal{A} and \mathcal{B} are *equal up to $p^n \cdot F^j$* (where $n, j \in \mathbf{Z}$; $n \geq 1$) if they are isomorphic modulo p^n , and, moreover, the difference between the two deformations $\mathcal{A}_{\mathbf{Z}/p^{n+1}\mathbf{Z}}$ and $\mathcal{B}_{\mathbf{Z}/p^{n+1}\mathbf{Z}}$ belongs to the F^j portion of the Hodge filtration on $\mathbf{R}^1 f_{\text{cr},*}(\text{Ad}(\mathcal{A}_{\mathbf{F}_p}))$.
- (2) Fix $i \in \mathbf{Z}$. Suppose that a deformation \mathcal{A}_i (respectively, \mathcal{A}_{i+1}) of $(\mathcal{P}_i)_{\mathbf{F}_p}$ (respectively, $(\mathcal{P}_{i+1})_{\mathbf{F}_p}$) up to $p^n \cdot F^0$ (respectively, modulo p^n) has been given. Then we shall refer to this pair of deformations as *suitable* if the relevant (i.e., naive or renormalized) Frobenius pull-back of \mathcal{A}_i is equal to \mathcal{A}_{i+1} (modulo p^n).

Now let us first observe that, by the above discussion, it follows that suitable deformations of the $(\mathcal{P}_i)_{\mathbf{F}_p}$ up to $p \cdot F^0$ exist, and are uniquely determined by the condition of “suitability” – cf. Proposition 1.5 of Chapter III (of the present work), as well as [Mzk1], Chapter II, Proposition 1.2; the discussion preceding Theorem 2.8 in [Mzk1], Chapter III.

Now we begin an inductive procedure (cf. the discussion preceding Theorem 2.8 of Chapter III of [Mzk1]). We will apply (ascending) induction on n to the following statement:

- ($*_n$) There exist unique suitable liftings of the $(\mathcal{P}_i)_{\mathbf{F}_p}$ up to $p^n \cdot F^0$.

Observe that ($*_1$) was just proven above.

Now let us fix n , and prove that ($*_{n+1}$) holds under the assumption that ($*_n$) holds. As a guide to visualizing what is going on, the reader may wish to use the Pictorial Appendix to this Chapter. We remark, however, that the descriptions there are coordinated so as to be compatible with the notation of the proof of Theorem 1.8. Thus, “ n ” here corresponds to “ $n-1$ ” in the Pictorial Appendix.

One checks easily – by considering global deformations of \mathcal{P}_i as the result of gluing together local deformations in the Zariski topology of

X , and using the fact that $F^0(\mathrm{Ad}(\mathcal{P}_i)_{\mathbf{F}_p})$ is a *Lie subalgebra* of $\mathrm{Ad}(\mathcal{P}_i)_{\mathbf{F}_p}$ – that:

The set of deformations of \mathcal{P}_i modulo $p^n \cdot F^0$ to a bundle modulo $p^{n+1} \cdot F^0$ naturally forms a torsor $\mathcal{T}_{n,i}$ over the vector bundle $H'_i \stackrel{\mathrm{def}}{=} \mathcal{H}'(\mathcal{P}_i)_{\mathbf{F}_p}$.

Pushing this torsor forward by means of the “change of structure group” morphism $H'_i \rightarrow \Theta'_i$ gives a torsor $\mathcal{R}_{n,i}$ over Θ'_i . We will say that two deformations are *equal modulo* $p^{n+1} \cdot F^0 + p^n \cdot F^1$ if their images in $\mathcal{R}_{n,i}$ coincide.

Let us first consider the case $i = -1$. By the induction hypothesis on n , \mathcal{P}_{-1} and \mathcal{P}_0 are determined up to $p^n \cdot F^0$. On the other hand, by the discussion above on infinitesimal deformations, the fact that Φ_{-1} is an isomorphism implies that there exists a unique deformation of \mathcal{P}_{-1} up to $p^{n+1} \cdot F^0 + p^n \cdot F^1$ whose relevant Frobenius pull-back coincides with \mathcal{P}_0 up to $p^n \cdot F^1$ (which is the same as $p^n \cdot F^0$ since $\Pi(0) = \chi$). Thus, in particular, \mathcal{P}_{-1} is determined up to $p^{n+1} \cdot F^0 + p^n \cdot F^1$. Repeating this argument (this time using the fact that Φ_{-2} is an isomorphism) thus shows that a \mathcal{P}_{-2} whose relevant Frobenius pull-back equals \mathcal{P}_{-1} up to $p^n \cdot F^1$ exists and is unique up to $p^{n+1} \cdot F^0 + p^n \cdot F^1$. Continuing in this fashion (using descending induction on i and the fact that Φ_i is an isomorphism), shows that, for all i , there exists a unique \mathcal{P}_i up to $p^{n+1} \cdot F^0 + p^n \cdot F^1$ whose relevant Frobenius pull-back is equal to \mathcal{P}_{i+1} up to $p^n \cdot F^1$. Now, for any i , the relevant Frobenius pull-back of \mathcal{P}_{i-1} is defined modulo p^{n+1} , so we obtain liftings of \mathcal{P}_i modulo p^{n+1} . But we also have (compatible) liftings of \mathcal{P}_i up to $p^{n+1} \cdot F^0 + p^n \cdot F^1$. Thus, in summary, we see that we have constructed a unique system of suitable \mathcal{P}_i up to $p^{n+1} \cdot F^0$. In other words, we have proven (\ast_{n+1}) . This completes the induction on n , and hence the proof of the Lemma.

○

Let us review our situation. On the one hand, we have a morphism $\mathcal{N}^{\mathrm{s},\mathrm{ord},\mathrm{pro}} \rightarrow \mathcal{Q}^{\mathrm{pro}}$ by the preceding Lemma. Clearly, this morphism factors through $\mathcal{Q}^{\mathrm{pro},\mathrm{ord}}$, so we have a morphism

$$\alpha : \mathcal{N}^{\mathrm{s},\mathrm{ord},\mathrm{pro}} \rightarrow \mathcal{Q}^{\mathrm{pro},\mathrm{ord}}$$

On the other hand, $\mathcal{Q}^{\mathrm{pro},\mathrm{ord}} \rightarrow \mathcal{N}^{\mathrm{s},\mathrm{ord}}$ defines (since $\mathcal{Q}^{\mathrm{pro},\mathrm{ord}}$ is perfect) a morphism

$$\beta : \mathcal{Q}^{\mathrm{pro},\mathrm{ord}} \rightarrow \mathcal{N}^{\mathrm{s},\mathrm{ord},\mathrm{pro}}$$

Clearly, $\beta \circ \alpha$ is the identity on $\mathcal{N}^{\mathrm{s},\mathrm{ord},\mathrm{pro}}$. On the other hand, it follows easily from the uniqueness part of Lemma 1.5 that $\alpha \circ \beta$ must also be the identity. Thus, in summary, we have the following result:

Theorem 1.6. *The natural morphism $\mathcal{Q}^{\text{pro,ord}} \rightarrow \mathcal{N}^{\text{s,ord,pro}}$ is an isomorphism.*

This gives us a very explicit description of the perfect log algebraic stack $\mathcal{Q}^{\text{pro,ord}}$.

§1.3. Systems of Canonical Modular Frobenius Liftings

Let $S_{\mathbf{F}_p}^{\log} = ((\overline{\mathcal{N}}_{g,r}^{\Pi,s})^{\log})^{\text{ord}}$. Thus, $S_{\mathbf{F}_p}^{\log}$ is étale (even without the log structures) and of finite type over $(\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{F}_p}$. It thus follows that there exists a unique p -adic formal stack

$$((\overline{\mathcal{N}}_{g,r}^{\Pi,s})^{\log})_{\mathbf{Z}_p}^{\text{ord}}$$

(which, here, for the sake of convenience, we shall denote by S^{\log}) which is étale over $(\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{Z}_p}$. In this subsection, we shall observe that the lifting procedure carried out in Lemma 1.5 can be modified slightly so as to give rise to certain *canonical modular Frobenius liftings* on S^{\log} . Note that this is in the spirit of Theorem 2.8 of Chapter III of [Mzk1]. Since the construction is very similar to that of Lemma 1.5 above and Theorem 2.8 of Chapter III of [Mzk1], we will omit many obvious details.

The system of Frobenius liftings and bundles that we obtain will be of the following nature: First note that from the structure morphism $S^{\log} \rightarrow \overline{\mathcal{M}}_{g,r}^{\log}$, we obtain a tautological log-curve $X^{\log} \rightarrow S^{\log}$. If $i \in \mathbf{Z}$ is *indigenous* (i.e., $\Pi(i) = \chi$), then we will obtain an *indigenous bundle*

$$\mathcal{P}_i$$

on X^{\log} whose reduction modulo p is the tautological (i.e., given by the definition of $(S^{\log})_{\mathbf{F}_p}$) i^{th} link bundle. Suppose that $i, j \in \mathbf{Z}$ are such that (i, j) are Π -ind-adjacent (i.e., $\Pi(i) = \Pi(j) = \chi$, and $\Pi(m) \neq \chi$ for every integer m such that $i < m < j$). Then we will obtain a *Frobenius lifting*

$$\Phi_j^{\log} : S^{\log} \rightarrow S^{\log}$$

i.e., a morphism of log schemes whose reduction modulo p will be equal to the $(j - i)^{\text{th}}$ power of the Frobenius morphism on $(S^{\log})_{\mathbf{F}_p}$.

Definition 1.7. The above Frobenius liftings and bundles will be said to be *compatible* if the following are satisfied:

- (1) First, note that for each $l \in \mathbf{Z}$ such that $0 \leq l < j - i$, we can form a crystal

$$\mathcal{P}_i[l] \text{ on } \text{Crys}((X^{\log})_{\mathbf{F}_p}^{F^l}/S^{\log})$$

as follows: If we pull-back \mathcal{P}_i via Φ_j^{\log} , then we get a crystal

$$(\Phi_j^{\log})^* \mathcal{P}_i \text{ on } \text{Crys}((X^{\log})_{\mathbf{F}_p}^{F^{j-i}}/S^{\log})$$

By forming the *renormalized Frobenius pull-back* of this crystal, we get a crystal

$$\mathcal{P}_i[j - i - 1] \text{ on } \text{Crys}((X^{\log})_{\mathbf{F}_p}^{F^{j-i-1}}/S^{\log})$$

Then the first part of the *compatibility* states that $\mathcal{P}_i[j - i - 1]_{\mathbf{F}_p}$ be equal to the tautological $(i + 1)^{\text{th}}$ link bundle on $(X^{\log})_{\mathbf{F}_p}^{F^{j-i-1}}$. Continuing in this fashion, that is to say, by executing a naive/renormalized Frobenius pull-back of $\mathcal{P}_i[l]$ to form

$$\mathcal{P}_i[l - 1] \text{ on } \text{Crys}((X^{\log})_{\mathbf{F}_p}^{F^{l-1}}/S^{\log})$$

and then stipulating (as part of the *compatibility*) that $\mathcal{P}_i[l - 1]_{\mathbf{F}_p}$ be equal to the tautological $(j - l + 1)^{\text{th}}$ link bundle on $(X^{\log})_{\mathbf{F}_p}^{F^{l-1}}$, we see that we obtain $\mathcal{P}_i[l]$ for all l such that $0 \leq l < j - i$.

- (2) The second part of the *compatibility* is that $\mathcal{P}_i[0] = \mathcal{P}_j$.

Moreover, in the following, we will also often consider *compatibility* modulo various powers of p or (more generally) modulo such symbols as $p^n \cdot F^m$ (as explained in the proof of Lemma 1.5).

Now, we start the actual construction. Just as in the proof of Lemma 1.5, we will use the terminology “equal up to $p^n \cdot F^m$,” etc., for comparing two crystals whose reductions modulo p are crys-stable bundles. For a visual aid to understanding what is going on in the following lifting procedure, we refer the reader to the Pictorial Appendix at the end of this Chapter.

Let us begin the *first step*. First of all, modulo p , all the $\mathcal{P}_i[l]$ ’s are completely determined by the tautological link bundles (cf. Definition 1.7, (1)). Also, all the \mathcal{P}_i ’s are determined up to $p^2 \cdot F^{-1} + p \cdot F^0$ by the fact that they are indigenous. Moreover, by the relationship between Frobenius liftings and “FL-bundles” (cf. [Mzk1], Chapter II,

§1, especially Proposition 1.2), one sees immediately that for each Π -ind-adjacent (i, j) , Φ_j^{\log} is determined modulo p^2 by the condition that $\mathcal{P}_i[j - i - 1]_{\mathbf{F}_p}$ be the appropriate tautological link bundle. By (ascending) induction on l (where $0 \leq l < j - i$), one sees that the $\mathcal{P}_i[l]$ are also all determined up to $p^2 \cdot F^{-1} + p \cdot F^0$ (cf. Chapter III, Proposition 1.5; [Mzk1], Chapter II, Proposition 1.2).

Now we begin the n^{th} step (where $n \geq 2$). The *hypothesis* for the n^{th} step is that the following objects have been determined, and are unique subject to the condition of *compatibility* (as defined above) modulo as high a power of p as is possible:

- (1) the Frobenius liftings Φ_j^{\log} modulo p^n ;
- (2) the crystals $\mathcal{P}_i[l]$ modulo $p^n \cdot F^{-1} + p^{n-1} \cdot F^0$.

Moreover, the \mathcal{P}_i are determined modulo $p^{n+1} \cdot F^{-1} + p^{n-1} \cdot F^0$ by the fact that they are to be indigenous.

Now (ascending) *induction* on l plus the definition of “ Π -ordinary” (cf. Definition 1.1) show that the $\mathcal{P}_i[l]$ are determined modulo $p^{n+1} \cdot F^{-1} + p^n \cdot F^0 + p^{n-1} \cdot F^1$. In particular, it follows that $\mathcal{P}_i[j - i + 1]$ is determined modulo $p^n \cdot F^{-1} + p^{n-1} \cdot F^1$. It thus follows from Π -ordinariness that Φ_j^{\log} is determined modulo p^{n+1} by the condition that we get the right $\mathcal{P}_i[j - i + 1]$ modulo $p^n \cdot F^{-1} + p^{n-1} \cdot F^1$.

Finally, by *executing the appropriate Frobenius pull-back* of \mathcal{P}_i or $\mathcal{P}_i[l]$, we see that all the $\mathcal{P}_i[l]$ are also determined up to $p^n \cdot F^{-1}$. Thus, in summary, it follows that the $\mathcal{P}_i[l]$ are, in fact, determined up to $p^{n+1} \cdot F^{-1} + p^n \cdot F^0$. That is to say, all the hypotheses necessary for the $(n + 1)^{\text{th}}$ step are satisfied. This completes the induction, and hence the proof of the following:

Theorem 1.8. *Let $S^{\log} = ((\overline{\mathcal{N}}_{g,r}^{\Pi,s})^{\log})_{\mathbf{Z}_p}^{\text{ord}}$. Let $X^{\log} \rightarrow S^{\log}$ be the tautological log-curve pull-back from the étale structure morphism $S^{\log} \rightarrow (\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{Z}_p}$. Then there exists a unique compatible system consisting of*

- (1) *for each indigenous i , an indigenous bundle \mathcal{P}_i on X^{\log}*
- (2) *for each Π -ind-adjacent (i, j) , a Frobenius lifting $\Phi_j^{\log} : S^{\log} \rightarrow S^{\log}$ (i.e., a morphism whose reduction modulo p is equal to the $(j - i)^{\text{th}}$ power of the Frobenius morphism on $(S^{\log})_{\mathbf{F}_p}$).*

We shall refer to this system of Frobenius liftings as the canonical system of modular Frobenius liftings on S^{\log} associated to the VF-pattern Π of period ϖ .

Finally, when (Π, ϖ) is the home VF-pattern, the resulting Frobenius lifting and indigenous bundle are the same as those of Theorem 2.8 of Chapter III of [Mzk1].

Before continuing, we would like to discuss the relationship between Theorems 1.6 and 1.8. Let $j_0 = 0 < j_1 < \dots < j_{s-1} < j_s = \varpi$ be the set of indigenous integers in a period of Π . Let

$$\Phi_{\text{tot}}^{\log} : S^{\log} \rightarrow S^{\log}$$

be the composite morphism $\Phi_{j_1}^{\log} \circ \dots \circ \Phi_{j_s}^{\log}$. Consider the inverse system

$$\dots \quad S^{\log} \quad \xrightarrow{\Phi_{\text{tot}}^{\log}} \quad S^{\log} \quad \dots$$

indexed by \mathbb{N} , all of whose objects are S^{\log} , and all of whose morphisms are Φ_{tot}^{\log} . Then note that the p -adic completion of the inverse limit of this inverse system is naturally isomorphic to $W((S^{\log})_{\mathbb{F}_p}^{\text{pro}})$. Thus, we shall identify the two in what follows. Another way to put this is that we have a commutative diagram

$$\begin{array}{ccc} W((S^{\log})_{\mathbb{F}_p}^{\text{pro}}) & \xrightarrow{\alpha} & S^{\log} \\ \downarrow & & \downarrow \Phi_{\text{tot}}^{\log} \\ W((S^{\log})_{\mathbb{F}_p}^{\text{pro}}) & \xrightarrow{\alpha} & S^{\log} \end{array}$$

where α is the projection to the first factor in the inverse limit, and the vertical arrow on the left is the ϖ^{th} power of the canonical Frobenius on $W((S^{\log})_{\mathbb{F}_p}^{\text{pro}})$. Let $Z^{\log} \rightarrow (S^{\log})_{\mathbb{F}_p}^{\text{pro}}$ be the pull-back of $X^{\log} \rightarrow S^{\log}$. Then, by pulling back by α , the indigenous bundle \mathcal{P}_0 on X^{\log} defines a crystal

$$\alpha^* \mathcal{P}_0$$

on $\text{Crys}(Z^{\log}/W((S^{\log})_{\mathbb{F}_p}^{\text{pro}}))$. On the other hand, by Theorem 1.6, $(S^{\log})_{\mathbb{F}_p}^{\text{pro}}$ is naturally isomorphic to $\mathcal{Q}^{\text{pro,ord}}$. We are now ready to state the relationship between Theorems 1.6 and 1.8:

Theorem 1.9. *Up to identifying $(S^{\log})_{\mathbb{F}_p}^{\text{pro}}$ and $\mathcal{Q}^{\text{pro,ord}}$ via the natural isomorphism of Theorem 1.6, the pulled-back crystal $\alpha^* \mathcal{P}_0$ is the tautological Π -indigenous crystal on $\text{Crys}(Z^{\log}/W(\mathcal{Q}^{\text{pro,ord}}))$ arising from the definition of $\mathcal{Q}^{\text{pro,ord}}$.*

Proof. This follows formally from the uniqueness portion of Lemma 1.5. \circ

Remark on Functoriality Properties. Naturally, the objects (isomorphisms, Frobenius liftings, crystals, etc.) constructed in Theorems 1.6, 1.8, and 1.9 are all compatible with the clutching morphisms that map

products of $\overline{\mathcal{M}}_{g_i, r_i}$'s into the boundary of $\overline{\mathcal{M}}_{g, r}$, as well as with log admissible coverings (cf. Theorems 2.8, 2.10 of Chapter III of [Mzk1]). We leave to the reader the tedious task of writing out and proving these various compatibilities. The proofs are all trivial.

§1.4. The Case of Elliptic Curves

Before continuing, we note that the theory of Chapters III, VI and VII (which was developed under the assumption of hyperbolicity, i.e., $2g - 2 + r \geq 1$) has an analogue in the case of elliptic curves, i.e., when $g = 1$, and $r = 0$. We will sketch here briefly what happens for elliptic curves; the proofs and formal definitions are all similar to the hyperbolic case, only technically much simpler.

First, we begin with *VF-patterns*. Since $\chi = 2g - 2 = 0$, it is easy to see that it only makes sense to consider two types of bundles: “indigenous” and “non-indigenous.” That is to say, unlike the hyperbolic case, there are no intermediate levels between “ χ ” and “0.” Thus, for elliptic curves, we define the set of levels (cf. Chapter III, §1.1), $\mathcal{L}ev$, to be the set consisting of the two symbols “ χ ” and “0,” i.e.,

$$\mathcal{L}ev = \{0, \chi\}$$

Now, I *claim* that

For *ordinary* elliptic curves, “it only makes sense” to consider the VF-pattern such that $\varpi = 1$ and $\Pi(n) = \chi$ for all $n \in \mathbf{Z}$. We shall refer to this VF-pattern below as the *ordinary pattern*.

Indeed, if there existed an $n \in \mathbf{Z}$ such that $\Pi(n) = 0$, $\Pi(n+1) = \chi$, then this would correspond to the situation in which the link bundle numbered $n+1$ is a *dormant* indigenous bundle. But by Proposition 2.4, (3), of Chapter IV, this implies that the m^{th} link bundle for all $m \leq n$ must also be *dormant*, i.e., that $\Pi(m) = 0$ for all $m \leq n$. In other words, such a VF-pattern would not admit a finite period, hence would not correspond to indigenous bundles that are invariant under a finite number of applications of Frobenius. This shows that it only makes sense to consider Π such that $\Pi(n) = \chi$ for all $n \in \mathbf{Z}$. Moreover, the VF-stack in this case is easily seen to be the fibered product of ϖ copies of $\overline{\mathcal{N}}_{1,0}^{\text{ord}} = \overline{\mathcal{M}}_{1,0}^{\text{ord}}$ over $\overline{\mathcal{M}}_{1,0}$. But since the natural morphism from $\overline{\mathcal{N}}_{1,0}^{\text{ord}}$ to $\overline{\mathcal{M}}_{1,0}$ is an *open immersion*, such a fibered product is always equal to $\overline{\mathcal{N}}_{1,0}^{\text{ord}}$ itself. Thus, consideration of the case $\varpi > 1$ does not give us anything “new,” i.e., anything that is not already “contained” in the case $\varpi = 1$. This completes our verification of the above “claim.”

Next, I *claim* that

For *supersingular* elliptic curves, “it only makes sense” to consider those VF-patterns such that ϖ is *even* and $\Pi(2n) = \chi$, $\Pi(2n + 1) = 0$ for all $n \in \mathbf{Z}$. We shall refer to such a VF-pattern below as the *supersingular pattern of period ϖ* .

Indeed, by Proposition 2.5, (1), of Chapter IV, every nilpotent indigenous bundle on a supersingular elliptic curve is *dormant*. Thus, we necessarily get $\Pi(0) = \chi$, $\Pi(-1) = 0$. Moreover, by Proposition 2.5, (2), of Chapter IV, *no* nilpotent 1-connection is dormant. Thus, $\Pi(-2) \neq 0$, i.e., $\Pi(-2) = \chi$. Repeating this argument completes the verification of the claim.

Next, let us consider the theory of canonical p -adic liftings. The only VF-patterns to consider are the *ordinary pattern* and the *supersingular pattern of period ϖ* (for even ϖ). First of all:

In the case of the *ordinary pattern*, $\mathcal{Q} = \mathcal{Q}^{\text{pro}} = \mathcal{Q}^{\text{pro,ord}}$ is equal to “pro” of the ordinary locus $\overline{\mathcal{M}}_{1,0}^{\text{ord}}$ of $(\overline{\mathcal{M}}_{1,0})_{\mathbf{F}_p}$. In particular, \mathcal{Q}^{pro} is affine. Moreover, \mathcal{Q}^{pro} embeds as a closed substack of \mathcal{S}_W and \mathcal{M}_W (defined analogously to the hyperbolic case – cf. Chapter VI, §2.3).

That is to say, just as was observed in the final subsection of [Mzk1], Chapter III, §3, the theory of liftings of ordinary elliptic curves is entirely analogous to the theory of liftings of “classical ordinary” (i.e., ordinary in the sense of the theory of [Mzk1]) hyperbolic curves.

On the other hand, in the supersingular case, it is not difficult to see that for $S = \text{Spec}(k)$ (where k is an algebraically closed field of characteristic p), the S -valued points of \mathcal{Q} correspond to diagrams (cf. the discussion at the end of §1.5 of the Introduction)

$$E_0 \xrightarrow{\phi_0} E_1 \xrightarrow{\phi_1} \dots \xrightarrow{\phi_{n-2}} E_{n-1} \xrightarrow{\phi_{n-1}} E_n = E_0$$

where $2n = \varpi$; $E_0, E_1, \dots, E_n = E_0$ are elliptic curves over $W(k)$ that are all equal to the same supersingular elliptic curve E_S modulo p . Thus, E_S will, in fact, be defined over \mathbf{F}_{p^2} . Moreover, the ϕ_i ’s are all Φ_W^2 -linear (where Φ_W denotes the Frobenius morphism on $W(k)$) isogenies that are equal to the square of the Frobenius morphism modulo p . Put another way, in the language of Chapter VIII:

The data parametrized by points of $\mathcal{Q}(k)$ (for the *supersingular pattern of period ϖ*) is the data for (a degenerate version of what in the discussion of the hyperbolic case is called) an *anabelian geometry* of length $\frac{1}{2}\varpi$ and order 2 on an elliptic curve over $W(k)$ whose reduction modulo p is supersingular.

(cf. Chapter VIII, §2.5, §2.6, §3.2). Of course, the term “anabelian” (which was chosen for its suitability in the hyperbolic case) sounds somewhat inappropriate in the present discussion of elliptic curves. The “anabelian geometry” that appears in the present discussion is *degenerate* (relative to what happens in general in the hyperbolic case – cf. §2.6 of Chapter VIII) – in the sense that all the Frobenius liftings involved (i.e., the ϕ_i) *commute* with one another.

Moreover, it is easy to see (for instance by considering Dieudonné modules of the E_i modulo p) that each $\phi_i : E_i \rightarrow E_{i+1}$ is isomorphic to the morphism $E_i \rightarrow E_i^{F^2}$ given by composing the ($W(k)$ -linear) morphism $E_i \rightarrow E_i$ given by multiplication by p with the Φ_W^2 -linear isomorphism $E_i \rightarrow E_i^{F^2}$. In particular, we get a naturally induced $W(k)$ -linear isomorphism

$$E_{i+1} \cong E_i^{F^2}$$

In particular, by composing these isomorphisms over a period, we get a natural $W(k)$ -linear isomorphism

$$E_n \cong E_0^{F^\varpi}$$

i.e., a $W(\mathbf{F}_{p^\varpi})$ -structure on E_0 . Conversely, it is a formal consequence of this discussion that once one chooses a $W(\mathbf{F}_{p^\varpi})$ -structure on E_0 , the diagram $E_0 \rightarrow E_1 \rightarrow \dots \rightarrow E_n$ is uniquely determined, up to natural isomorphism. In summary, we have the following:

In the case of the *supersingular pattern of period ϖ* , \mathcal{Q}^{pro} is étale over the supersingular locus of $(\overline{\mathcal{M}}_{1,0})_{\mathbf{F}_p}$, and embeds as a closed substack of \mathcal{S}_W . Moreover, if k is an algebraically closed field of characteristic p , then the S -valued points of \mathcal{Q} correspond to elliptic curves over $W(\mathbf{F}_{p^\varpi})$ whose reductions modulo p are supersingular. In particular, $\mathcal{Q}^{\text{pro,ord}}$ is empty. Finally, if $\varpi = 2$, then $\mathcal{Q} = \mathcal{Q}^{\text{pro}}$ (cf. Chapter VI, Proposition 2.1).

Thus, it follows in particular that the image of the ordinary \mathcal{Q}^{pro} in \mathcal{S}_W is disjoint from the image of any of the supersingular $(\mathcal{Q}')^{\text{pro}}$ in \mathcal{S}_W (since the respective images in $(\overline{\mathcal{M}}_{1,0})_{\mathbf{F}_p}$ are disjoint).

§2. The Closure of the Binary Ordinary Locus

§2.1. The Deperfection of the Closure

Let Π be a *binary* (i.e., $\Pi(\mathbf{Z}) \subseteq \{0, \chi\}$) VF-pattern of period ϖ . Let \mathcal{Q}^{pro} be the (“pro” of the) corresponding stack of quasi-analytic self-isogenies. Recall (Chapter VI, Theorem 2.2) that \mathcal{Q}^{pro} is representable by a perfect algebraic log stack. Inside \mathcal{Q}^{pro} , we have an *open* substack

$$\mathcal{Q}^{\text{pro,ord}} \subseteq \mathcal{Q}^{\text{pro}}$$

given by the Π -ordinary locus. The purpose of the present section is to study the closure of $\mathcal{Q}^{\text{pro,ord}}$ in \mathcal{Q}^{pro} .

We begin by introducing some notation. Let us write

$$\mathcal{C} \subseteq \mathcal{Q}^{\text{pro}}$$

for the (schematic) closure of $\mathcal{Q}^{\text{pro,ord}}$ in $\mathcal{Q}^{\text{pro,ord}}$. Thus, \mathcal{C} is a perfect algebraic log stack over \mathbf{F}_p which admits $\mathcal{Q}^{\text{pro,ord}}$ as a dense open substack. Next, let us write

$$\mathcal{N} \stackrel{\text{def}}{=} ((\overline{\mathcal{N}}_{g,r}^{\Pi,s})^{\log})^{\text{ord}}$$

Thus, we may construct an inverse system, indexed by the *indigenous integers* (i.e., $i \in \mathbf{Z}$ such that $\Pi(i) = \chi$), as follows: For each such i , we take the object numbered “ i ” of this inverse system to be \mathcal{N} . Moreover, if $i < j$ are *ind-adjacent*, we take the morphism $\mathcal{N} \rightarrow \mathcal{N}$ numbered j (in the inverse system) to be the canonical Frobenius lifting $\Phi_j^{\log} : \mathcal{N} \rightarrow \mathcal{N}$ of Theorem 1.8. (Here, for $j < 0$ or $j > \varpi$, we define Φ_j^{\log} by the condition that the function $j \mapsto \Phi_j^{\log}$ be periodic of period ϖ .) Thus, we have an inverse system

$$\dots \xrightarrow{\Phi_2^{\log}} \mathcal{N} \xrightarrow{\Phi_1^{\log}} \mathcal{N} \xrightarrow{\Phi_0^{\log}} \mathcal{N} \xrightarrow{\Phi_{-1}^{\log}} \mathcal{N} \xrightarrow{\Phi_{-2}^{\log}} \dots$$

Moreover, by Theorems 1.6 and 1.9, we know that the p -adic completion of the inverse limit of this system may be naturally identified with $W(\mathcal{Q}^{\text{pro,ord}})$. Write

$$\alpha_i : W(\mathcal{Q}^{\text{pro,ord}}) \rightarrow \mathcal{N}$$

for the composite of $\Phi_{W(\mathcal{Q}^{\text{pro,ord}})}^i$ (i.e., the result of iterating the natural Frobenius on $W(\mathcal{Q}^{\text{pro,ord}})$ a total of i times) with the projection of the

inverse limit $W(\mathcal{Q}^{\text{pro,ord}})$ to the i^{th} member of the inverse system. Thus, it is a tautology that we have a commutative diagram

$$\begin{array}{ccc} W(\mathcal{Q}^{\text{pro,ord}}) & \xrightarrow{\alpha_j} & \mathcal{N} \\ \downarrow \Phi_{W(\mathcal{Q}^{\text{pro,ord}})}^{j-i} & & \downarrow \Phi_j^{\text{log}} \\ W(\mathcal{Q}^{\text{pro,ord}}) & \xrightarrow{\alpha_i} & \mathcal{N} \end{array}$$

Moreover, note that the function $i \mapsto \{\alpha_i : W(\mathcal{Q}^{\text{pro,ord}}) \rightarrow \mathcal{N}\}$ is periodic of period ϖ (cf. the discussion preceding Theorem 1.9), while the $(\alpha_i)_{\mathbf{F}_p}$'s are all the same (i.e., independent of i).

Next, observe that we have natural inclusions $W(\mathcal{O}_{\mathcal{C}}) \hookrightarrow W(\mathcal{O}_{\mathcal{Q}^{\text{pro,ord}}})$ (induced by $\mathcal{Q}^{\text{pro,ord}} \subseteq \mathcal{C}$) and $\alpha_i^{-1} : \mathcal{O}_{\mathcal{N}} \hookrightarrow W(\mathcal{O}_{\mathcal{Q}^{\text{pro,ord}}})$ (induced by $\alpha_i : W(\mathcal{Q}^{\text{pro,ord}}) \rightarrow \mathcal{N}$). Let us denote by $\mathcal{O}_{\mathcal{D}_i}$ the sheaf of rings on the étale site of (the non-log algebraic stack underlying) \mathcal{C} obtained by taking the *intersection* of $W(\mathcal{O}_{\mathcal{C}})$ with the image of α_i^{-1} (inside $W(\mathcal{O}_{\mathcal{Q}^{\text{pro,ord}}})$). Thus, we have a commutative diagram:

$$\begin{array}{ccc} \mathcal{O}_{\mathcal{D}_i} & \longrightarrow & W(\mathcal{O}_{\mathcal{C}}) \\ \downarrow & & \downarrow \\ \mathcal{O}_{\mathcal{N}} & \longrightarrow & W(\mathcal{O}_{\mathcal{Q}^{\text{pro,ord}}}) \end{array}$$

Write \mathcal{D}_i for the geometric object associated to $\mathcal{O}_{\mathcal{D}_i}$. Thus, \mathcal{D}_i is a p -adic formal algebraic log stack with the property that \mathcal{C} may be identified with $(\mathcal{D}_i)_{\mathbf{F}_p}^{\text{pro}}$. (Here, the log structure is obtained by taking an intersection of sheaves of étale monoids similar to the intersection used to define the structure sheaf $\mathcal{O}_{\mathcal{D}_i}$.) That is to say, \mathcal{D}_i is a sort of “de-perfection” of $W(\mathcal{C})$. Note that the function $i \mapsto \mathcal{D}_i$ is *periodic of period* ϖ .

Proposition 2.1. *The p -adic formal algebraic log stack \mathcal{D}_i is \mathbf{Z}_p -flat, contains \mathcal{N} as an open dense subobject, and satisfies $\mathcal{C} = (\mathcal{D}_i)_{\mathbf{F}_p}^{\text{pro}}$. Moreover, the natural morphisms $\mathcal{O}_{\mathcal{D}_i}/(p^n) \rightarrow \mathcal{O}_{\mathcal{N}}/(p^n)$, $\mathcal{O}_{\mathcal{D}_i}/(p^n) \rightarrow W(\mathcal{O}_{\mathcal{C}})/(p^n)$ (for $n \geq 0$) are injective, and $(\mathcal{D}_i)_{\mathbf{F}_p}$ is independent of i . Finally, the canonical Frobenius lifting $\Phi_j^{\text{log}} : \mathcal{N} \rightarrow \mathcal{N}$ extends naturally to a morphism (which, by abuse of notation, we write) $\Phi_j : \mathcal{D}_j \rightarrow \mathcal{D}_i$.*

Proof. This follows immediately from the definitions. Note that Φ_j^{log} extends to a morphism $\mathcal{D}_j \rightarrow \mathcal{D}_i$ by the commutative diagram above (that appears in the discussion following the definition of the α_i), plus the fact that the natural Frobenius on $W(\mathcal{Q}^{\text{pro,ord}})$ extends to the natural Frobenius on $W(\mathcal{C})$. \square

Definition 2.2. We shall refer to \mathcal{D}_i as the i^{th} *déperfection* of \mathcal{C} . Since $(\mathcal{D}_i)_{\mathbf{F}_p}$ is independent of i , we shall write \mathcal{D} for the common object $(\mathcal{D}_i)_{\mathbf{F}_p}$.

In the discussions of the following subsections, we will frequently need to consider the composite

$$\mathcal{D}_0 \xrightarrow{\Phi_0} \dots \xrightarrow{\Phi_j} \mathcal{D}_i$$

(where $j < i$ are ind-adjacent, and $i < 0$). Let us denote this composite by

$$\Psi_i : \mathcal{D}_0 \rightarrow \mathcal{D}_i$$

Next, let us write $\mathcal{M} \stackrel{\text{def}}{=} (\overline{\mathcal{M}}_{g,r})_{\mathbf{Z}_p}$ for the log stack of stable log-curves (of type (g, r)) over \mathbf{Z}_p ; $\mathcal{S} \stackrel{\text{def}}{=} (\overline{\mathcal{S}}_{g,r})_{\mathbf{Z}_p}$ for the log stack of stable log-curves (of type (g, r)) equipped with an indigenous bundle (over \mathbf{Z}_p). Then note that the indigenous bundle \mathcal{P}_i of Theorem 1.8 defines a classifying morphism $\mathcal{N} \rightarrow \mathcal{S}$ whose composite with $\alpha_i : W(\mathcal{Q}^{\text{pro,ord}}) \rightarrow \mathcal{S}$ is *equal* (cf. Theorem 1.9) to the composite of the classifying morphism $W(\mathcal{C}) \rightarrow \mathcal{S}$ (induced by the bundle \mathcal{P}_i of Chapter VI, §1.4) with the open immersion $W(\mathcal{Q}^{\text{pro,ord}}) \hookrightarrow W(\mathcal{C})$. Thus, by the definition of \mathcal{D}_i , we get a morphism

$$\pi_{\mathcal{S},i} : \mathcal{D}_i \rightarrow \mathcal{S}$$

(and, by composite with the projection $\mathcal{S} \rightarrow \mathcal{M}$) a morphism

$$\pi_{\mathcal{M},i} : \mathcal{D}_i \rightarrow \mathcal{M}$$

Note that $(\pi_{\mathcal{M},i})_{\mathbf{F}_p}$ is *independent of i* . In the following subsection, we would like to study *the derivative of $\pi_{\mathcal{M},i}$* .

§2.2. The Differentials of the Déperfection

We maintain the notation of the preceding subsection. Let us denote by

$$\Omega'_i$$

(for $i \in \mathbf{Z}$ indigenous for Π) the image of the sheaf of differentials $\Omega_{(\mathcal{D}_i)_{\mathbf{F}_p}}$ in (the push-forward to $(\mathcal{D}_i)_{\mathbf{F}_p}$ under the natural inclusion $\mathcal{N}_{\mathbf{F}_p} \hookrightarrow (\mathcal{D}_i)_{\mathbf{F}_p}$ of) $\Omega_{\mathcal{N}_{\mathbf{F}_p}}$. Thus, Ω'_i is a quasi-coherent sheaf of $\mathcal{O}_{(\mathcal{D}_i)_{\mathbf{F}_p}}$ -modules which

is a quotient of $\Omega_{(\mathcal{D}_i)_{\mathbf{F}_p}}$. Moreover, Ω'_i corresponds to Ω'_j under the natural identification of $(\mathcal{D}_i)_{\mathbf{F}_p}$ with $(\mathcal{D}_j)_{\mathbf{F}_p}$. Thus, we shall write Ω' for the corresponding sheaf on \mathcal{D} .

Next, let us observe that the derivative of $(\Phi_i)_{\mathbf{F}_p}$ is zero (since $(\Phi_i)_{\mathbf{F}_p}$ is a power of the Frobenius morphism). Moreover, by dividing the derivative of Φ_j by p , we get a morphism

$$\frac{1}{p} \cdot d\Phi_j : \Phi_j^* \Omega'_i \rightarrow \Omega'_j$$

(where $i < j \leq 0$; i and j are ind-adjacent). Indeed, this is not entirely obvious since the sheaf of (p -adically continuous) differentials $\Omega_{\mathcal{D}_i/\mathbf{Z}_p}$ need not be \mathbf{Z}_p -flat (although $\Omega_{\mathcal{N}/\mathbf{Z}_p}$ is, of course, \mathbf{Z}_p -flat – since \mathcal{N} is formally smooth over \mathbf{Z}_p). Nevertheless, we still get a morphism as desired, since Ω'_i is generated by d of sections f of $\mathcal{O}_{\mathcal{D}_i}$, and $\Phi_j^{-1}(f)$ may be locally written in the form $p \cdot h_1 + h_2^p$ (where h_1, h_2 are local sections of $\mathcal{O}_{\mathcal{D}_j}$); moreover, $\frac{1}{p} \cdot d(p \cdot h_1 + h_2^p) = dh_1 + h_2^{p-1} \cdot dh_2$ manifestly defines (when reduced modulo p) a section of Ω'_j , as desired. By composing the various $\frac{1}{p} \cdot d\Phi_j$'s, we thus get a morphism

$$\frac{1}{p^n} \cdot d\Psi_i : (\Phi_{\mathcal{D}}^*)^{|i|} \Omega'_i \rightarrow \Omega'_0$$

where $\Phi_{\mathcal{D}}$ is the Frobenius morphism on $\mathcal{D} = (\mathcal{D}_i)_{\mathbf{F}_p} = (\mathcal{D}_0)_{\mathbf{F}_p}$, and n is the number of negative indigenous integers $\geq i$. Let us write

$$\mu_i \stackrel{\text{def}}{=} \left(\frac{1}{p^n} \cdot d\Psi_i \right) \circ \{ (\Phi_{\mathcal{D}}^*)^{|i|} (d\pi_{\mathcal{M},i}) \} : (\Phi_{\mathcal{D}}^*)^{|i|} (\Omega_{\mathcal{M}_{\mathbf{F}_p}}|_{\mathcal{D}}) \rightarrow \Omega'_0$$

where $\Omega_{\mathcal{M}_{\mathbf{F}_p}}|_{\mathcal{D}}$ denotes the pull-back of $\Omega_{\mathcal{M}_{\mathbf{F}_p}}$ to $(\mathcal{D}_i)_{\mathbf{F}_p} = (\mathcal{D}_0)_{\mathbf{F}_p}$ via $(\pi_{\mathcal{M}_i})_{\mathbf{F}_p} = (\pi_{\mathcal{M}_0})_{\mathbf{F}_p}$.

Next, if $j > i$ are ind-adjacent, let us write

$$\delta_{i,j} : \Omega_{\mathcal{M}_{\mathbf{F}_p}}|_{\mathcal{D}} \rightarrow (\Phi_{\mathcal{D}}^*)^{j-i} (\Omega_{\mathcal{M}_{\mathbf{F}_p}}|_{\mathcal{D}})$$

for the map obtained by composing the duals of the (various appropriate Frobenius conjugates of the) “ Φ_i ” of Definition 1.1, where we take the “ i ” of “ Φ_i ” to be $i, i+1, i+2, \dots, j-2, j-1$. Note here that although a priori, this composite is only defined over \mathcal{C} , it is clear from the definition of \mathcal{D}_i that it is in fact defined over $(\mathcal{D}_i)_{\mathbf{F}_p}$, as desired. Also, it is clear that in the present *binary* context, the relevant “ $\Theta(\mathcal{P}_i)$ ” and “ $\Theta'(\mathcal{P}_i)$ ” of the discussion preceding Definition 1.1 may be identified with the dual of $\Omega_{\mathcal{M}_{\mathbf{F}_p}}|_{\mathcal{D}}$ in the present discussion (cf. Chapter I, Proposition 4.5).

The following technical observation is the key result in the proof of the main theorems of this section:

Lemma 2.3. *If $j > i$ are ind-adjacent, then the following diagram commutes:*

$$\begin{array}{ccc} (\Phi_{\mathcal{D}}^*)^{j-i} \Omega_{\mathcal{M}_{\mathbf{F}_p}}|_{\mathcal{D}} & \longrightarrow & (\Phi_{\mathcal{D}}^*)^{j-i} \Omega'_i \\ \uparrow \delta_{i,j} & & \downarrow \frac{1}{p} d\Phi_j \\ \Omega_{\mathcal{M}_{\mathbf{F}_p}}|_{\mathcal{D}} & \longrightarrow & \Omega'_j \end{array}$$

Here, the upper horizontal arrow is the morphism $(\Phi_{\mathcal{D}}^*)^{j-i} (d\pi_{\mathcal{M},i})$, while the lower horizontal arrow is the morphism $d\pi_{\mathcal{M},j}$.

Proof. Recall that we have a tautological indigenous bundle \mathcal{P}_i on some curve over \mathcal{D}_i . Moreover, if we take the renormalized Frobenius pull-back of \mathcal{P}_i , followed by $(j - i - 1)$ naive Frobenius pull-backs, we obtain the tautological indigenous bundle \mathcal{P}_j (on some curve over \mathcal{D}_j). Now the horizontal morphisms of the diagram in question are defined, respectively, by the derivatives of the classifying morphisms of \mathcal{P}_i and \mathcal{P}_j . Moreover, the various Frobenius pull-backs that were executed to obtain \mathcal{P}_j from \mathcal{P}_i were carried out over the Frobenius lifting $\Phi_j : \mathcal{D}_j \rightarrow \mathcal{D}_i$ – whose derivative is the vertical morphism on the right. On the other hand, the vertical morphism on the left is defined by the Frobenius action induced on the F^{-1}/F^0 portions of the respective de Rham cohomologies of \mathcal{P}_i and \mathcal{P}_j , which may be regarded as parametrizing infinitesimal deformations of \mathcal{P}_i and \mathcal{P}_j . Thus, the commutativity of the diagram in question is a tautology obtained by “differentiating the statement ‘the Frobenius pull-back of \mathcal{P}_i is \mathcal{P}_j .’ ” \circ

Remark. Lemma 2.3 above is a sort of generalized version of [Mzk1], Chapter III, Proposition 2.3 and Lemma 2.6.

Now let us return to the morphism

$$\mu_i : (\Phi_{\mathcal{D}}^*)^{[i]} (\Omega_{\mathcal{M}_{\mathbf{F}_p}}|_{\mathcal{D}}) \rightarrow \Omega'_0$$

defined above. Note that (by restricting to the ordinary locus) it is clear that μ_i is injective. Let us write

$$\Omega_{\mathcal{M}}[i] \subseteq \Omega'_0$$

for the image of μ_i . By repeated application of Lemma 2.3, we thus obtain:

Lemma 2.4. *We have*

$$\Omega_{\mathcal{M}}[j] \subseteq \Omega_{\mathcal{M}}[i] \subseteq \Omega'_0$$

where i and j are any nonpositive indigenous (not necessarily ind-adjacent!) integers such that $i < j$.

Now let us write

$$\omega'$$

for the image of $\wedge^{3g-3+r} \Omega'$ in (the push-forward to \mathcal{D} under the natural inclusion $\mathcal{N}_{\mathbf{F}_p} \hookrightarrow \mathcal{D}$ of) $\wedge^{3g-3+r} \Omega'|_{\mathcal{N}_{\mathbf{F}_p}}$. (Note that $3g-3+r$ is the dimension of $\mathcal{N}_{\mathbf{F}_p}$.) Thus, ω' is a quasi-coherent sheaf on \mathcal{D} whose restriction to the ordinary locus (i.e., $\mathcal{N}_{\mathbf{F}_p} \subseteq \mathcal{D}$) is a line bundle. Moreover, inside ω' , we have sub- $\mathcal{O}_{\mathcal{D}}$ -modules

$$\omega_{\mathcal{M}}[j] \subseteq \omega_{\mathcal{M}}[i] \subseteq \omega'$$

(for $i < j$ as in Lemma 2.4) defined by $\Omega_{\mathcal{M}}[j]$ and $\Omega_{\mathcal{M}}[i]$, respectively. Note that (since $\Omega_{\mathcal{M}_{\mathbf{F}_p}}|_{\mathcal{D}}$ is a vector bundle on \mathcal{D}), the $\omega_{\mathcal{M}}[i]$ are *line bundles*. Let us denote by

$$\mathcal{Z}(i, j) \subseteq \mathcal{D}$$

the schematic zero locus of the inclusion $\omega_{\mathcal{M}}[j] \subseteq \omega_{\mathcal{M}}[i]$. Note that we always have

$$\mathcal{N}_{\mathbf{F}_p} \subseteq \mathcal{D} - \mathcal{Z}(i, j)$$

(by the definition of “ Π -ordinary” – cf. Definition 1.1). (In fact, if $j - i \geq \varpi$, then this inclusion is even an equality.) Since the ordinary locus (i.e., $\mathcal{N}_{\mathbf{F}_p}$) is dense in \mathcal{D} , it thus follows that the $\mathcal{Z}(i, j)$ are *Cartier divisors* in \mathcal{D} .

Now let us take $j = 0$, and i to be a negative multiple of the period ϖ . Then, we have (for m a positive integer)

$$\omega_{\mathcal{M}}[0] \subseteq \omega_{\mathcal{M}}[-m \cdot \varpi] \subseteq \omega'$$

Here, we think of $\omega_{\mathcal{M}}[0]$ as “fixed,” and $\omega_{\mathcal{M}}[-m \cdot \varpi]$ as varying as $m \rightarrow \infty$. Moreover, by “the determinant of Lemma 2.3,” it follows that the Cartier divisors introduced above (written additively) satisfy:

$$[\mathcal{Z}(-m \cdot \varpi, 0)] = (1 + p^\varpi + p^{2 \cdot \varpi} + \dots + p^{(m-1) \cdot \varpi}) \cdot [\mathcal{Z}(-\varpi, 0)]$$

That is to say, $[\mathcal{Z}(-m \cdot \varpi, 0)]$ is a multiple of $[\mathcal{Z}(-\varpi, 0)]$ by a positive number which goes to infinity as $m \rightarrow \infty$.

In other words, roughly speaking, we have shown the following: Let f be a local equation defining the complement of the ordinary locus in \mathcal{D} . If we knew that $f^{-1} \in \mathcal{O}_{\mathcal{D}}$, then *this would imply that the ordinary locus $\mathcal{N}_{\mathbb{F}_p}$ is already closed in \mathcal{D}* . Although we do not quite obtain that strong a result, what we have done above implies that *the $\mathcal{O}_{\mathcal{D}}$ -module ω' is automatically a module over $\mathcal{O}_{\mathcal{D}}[f^{-1}]$* . That is to say,

“As far as the differentials are concerned, it is as if the ordinary locus is already closed in \mathcal{D} .”

In the next subsection, we will discuss this conclusion in greater detail.

§2.3. The ω -Closedness of the Binary Ordinary Locus

Let k be a perfect field. Let X be a reduced k -scheme (not necessarily locally of finite type over k !). Let $U \subseteq X$ be a dense open subscheme which is smooth and locally of finite type over k . Assume that U is of constant dimension d , and that the complement $X - U$ is set-theoretically a *Cartier divisor* (i.e., locally defined by a single equation). Write $\iota : U \hookrightarrow X$ for the natural inclusion. Note that the assumption on the complement $X - U$ implies that the push-forward by ι of a quasi-coherent sheaf on U is again quasi-coherent on X (cf. the proof of [Harts2], Chapter II, §5, Proposition 5.8, (c)). Write

$$\omega'_X \subseteq \iota_* \omega_U$$

for the image of ω_X in $\iota_* \omega_U$, where $\omega_U = \wedge^d \Omega_U$. Thus, ω'_X is a quasi-coherent \mathcal{O}_X -module, and $\omega'_X|_U = \omega_U$.

Definition 2.5. We shall say that U is ω -closed in X if $\omega'_X = \iota_* \omega_U$ (as quasi-coherent \mathcal{O}_X -modules).

Remark. There are obvious “log” and “algebraic stack” generalizations of Definition 2.5 (as well as of Propositions 2.6, 2.7, and 2.8 below). We leave it to the reader to spell out the unenlightening details of these easy generalizations.

Proposition 2.6. *If U is closed in X , then it is ω -closed in X .*

Proof. Indeed, since U is dense in X , if it is closed in X , then we have $U = X$, so it is clearly ω -closed in X . \circ

Proposition 2.7. *Suppose that X is locally of finite type over k . Then if U is ω -closed in X , then it is closed in X .*

Proof. First note that by replacing X by its normalization (which clearly will not effect the ω -closedness or the closedness of U), we may assume that X is *normal*. Also, since the problem is clearly local on X , we may assume that the complement $X - U$ is set-theoretically defined by a single equation $f \in \Gamma(X, \mathcal{O}_X)$. But now observe that since $\iota_*\omega_U$ is an $\mathcal{O}_X[f^{-1}]$ -module, it follows that ω'_X is also a $\mathcal{O}_X[f^{-1}]$ -module. On the other hand, ω'_X is also clearly *coherent* on X . This implies that f^{-1} is *integral* over \mathcal{O}_X , hence (since X is normal) that $f^{-1} \in \Gamma(X, \mathcal{O}_X)$, i.e., that $U = X$, as desired. \circ

Proposition 2.8. *Suppose that $d = 1$. Then if U is ω -closed in X , then it is closed in X .*

Proof. As in the proof of Proposition 2.7, we may assume that X is *normal*. Also, since the problem is clearly local on X , we may assume that X is *affine*, i.e., $X = \text{Spec}(A)$. Thus, A may be written as the inductive limit of a system of normal, one-dimensional k -algebras of finite type, all of which have the same quotient field. But it is well-known that such a system is necessarily trivial, i.e., A itself is necessarily of finite type over k . (Indeed, this is a consequence of the uniqueness of the natural compactification of a smooth curve over k – cf., e.g., [Harts2], Chapter I, §6.) Thus, the result follows from Proposition 2.7. \circ

Example 2.9. Consider the case $X \stackrel{\text{def}}{=} \text{Spec}(A)$, where

$$A \stackrel{\text{def}}{=} k[x, y, \{y \cdot x^{-i}\}_{i \in \mathbf{Z}_{\geq 0}}]$$

(and x and y are indeterminates). Let $U \stackrel{\text{def}}{=} \text{Spec}(A[\frac{1}{x}])$. Then since $x^{-1} \notin A$, it follows that U is *not closed* in X . On the other hand,

$$dx \wedge (x^{-i} \cdot dy) \in \omega'_A$$

so it follows easily that U is ω -closed in X . Thus, in particular, if the hypotheses of Propositions 2.7 and 2.8 are not satisfied, then there exist counter-examples to the converse of Proposition 2.6.

Now let U be any k -smooth scheme of locally finite type (over k) and constant dimension d . Let us assume, moreover, that k is of characteristic p . Thus, we may form U^{pro} (cf. Chapter VI, Definition 1.9). Then U^{pro} will be a perfect k -scheme, which we may think of as the “perfection” of U . Assume, moreover, that Y is a perfect k -scheme such that we have an embedding

$$U^{\text{pro}} \subseteq Y$$

of U^{pro} as a (not necessarily open or closed) subscheme of Y . Let Z be the schematic closure of U^{pro} in Y . Thus, Z is a perfect k -scheme. Moreover, if $\iota^{\text{pro}} : U^{\text{pro}} \hookrightarrow Z$ is the natural inclusion, then we have an injection $\mathcal{O}_Z \hookrightarrow \iota_*^{\text{pro}} \mathcal{O}_{U^{\text{pro}}}$. On the other hand, we also have a natural morphism $\alpha : U^{\text{pro}} \rightarrow U$ which gives rise to a morphism $\alpha^{-1} : \mathcal{O}_U \rightarrow \mathcal{O}_{U^{\text{pro}}}$. Let \mathcal{O}_X be the sheaf of rings on Z given by taking the intersection of the image of α^{-1} with the image of \mathcal{O}_Z in $\iota_*^{\text{pro}} \mathcal{O}_{U^{\text{pro}}}$. Then, if we equip the underlying space of Z with \mathcal{O}_X , one sees easily that we obtain a k -scheme X (with structure sheaf \mathcal{O}_X). Moreover, we have an open, dense embedding $U \hookrightarrow X$, as well as a commutative diagram:

$$\begin{array}{ccc} U^{\text{pro}} & \longrightarrow & Z \\ \downarrow & & \downarrow \\ U & \longrightarrow & X \end{array}$$

Here, Z may be identified with X^{pro} . Then we make the following:

Definition 2.10. We shall say that U^{pro} is ω -closed in Y if (U and X satisfy the hypotheses discussed as the beginning of this section and) U is ω -closed in X (as in Definition 2.5).

Note that the notion “ U^{pro} is ω -closed in Y ” is only defined once one specifies the auxiliary structure of U^{pro} as the “pro” of U .

Remark. As before, there are obvious “log” and “algebraic stack” generalizations of Definition 2.10. We leave it to the reader to spell out the unenlightening details of these easy generalizations.

Now we are ready to state the main theorem of this section:

Theorem 2.11. Let Π be a binary VF-pattern of period ϖ . Then $\mathcal{Q}^{\text{pro,ord}}$ is both open and ω -closed in \mathcal{Q}^{pro} . In particular,

- (1) If $3g - 3 + r = 1$, then $\mathcal{Q}^{\text{pro,ord}}$ is actually closed in \mathcal{Q}^{pro} .

- (2) If $\mathcal{R} \subseteq \mathcal{Q}^{\text{pro}}$ is a subobject containing $\mathcal{Q}^{\text{pro,ord}}$ and which is “pro” (cf. Chapter VI, Definition 1.9) of a fine algebraic log stack which is locally of finite type over \mathbf{F}_p , then $\mathcal{Q}^{\text{pro,ord}}$ is closed in \mathcal{R} .

In other words, at least among perfection of fine algebraic log stacks which are locally of finite type over \mathbf{F}_p , $\mathcal{Q}^{\text{pro,ord}}$ is already “complete” inside \mathcal{Q}^{pro} .

Proof. This follows immediately from the discussion as the end of §2.2, as well as Propositions 2.7 and 2.8. \circ

Corollary 2.12. *Let Π be a VF-pattern of pure tone ϖ . Then the associated $\mathcal{Q}^{\text{pro,ord}}$ embeds naturally as an ω -closed substack of S_W .*

Proof. This follows immediately from Corollary 2.6 of Chapter VI and Theorem 2.11 above. \circ

Corollary 2.13. *Let Π be a pre-home VF-pattern of period ϖ . Then the associated $\mathcal{Q}^{\text{pro,ord}}$ embeds naturally as an ω -closed substack of both S_W and M_W . Moreover, the image of $\mathcal{Q}^{\text{pro,ord}}$ in S_W is disjoint from the closure of any $(\mathcal{Q}')^{\text{pro}}$ associated to a non-pre-home Π' .*

Proof. This follows immediately from Theorem 2.11 above, and Propositions 2.3 and 2.9, and Corollary 2.13 of Chapter VI. \circ

§3. Existence Results

§3.1. The Binary Case

By Theorem 2.9 of Chapter IV, if $p > 4^{3g-3}$, then the *Lubin-Tate stack*

$$\mathcal{L}_g^\varpi$$

(cf. Chapter IV, Definition 2.8) is nonempty. But \mathcal{L}_g^ϖ is clearly just the Π -ordinary locus of $\overline{\mathcal{N}}_g^{\Pi,s}$, where Π is the VF-pattern of pure tone ϖ .

Now suppose that Π is a *binary* VF-pattern of period ϖ . Then it is easy to see that the Π -ordinary locus of $\overline{\mathcal{N}}_{g,r}^{\Pi,s}$ is just equal to a fibered product (over $\overline{\mathcal{M}}_{g,r}$) of various pure tone $\overline{\mathcal{N}}_{g,r}^{\Pi,s,\text{ord}}$'s. Thus, we obtain the following existence result:

Theorem 3.1. *Suppose that Π is a binary VF-pattern of period ϖ . If Π is pre-home, then the Π -ordinary locus of $\overline{\mathcal{N}}_{g,r}^{\Pi,s}$ is always nonempty. If $p > 4^{3g-3}$ and $r = 0$, and Π is arbitrary (but binary), then the Π -ordinary locus of $\overline{\mathcal{N}}_{g,r}^{\Pi,s}$ is nonempty.*

Now let us take a look at the system of Frobenius liftings that one obtains in the binary case. Let S^{\log} and $\Phi_j^{\log} : S^{\log} \rightarrow S^{\log}$ be as in Theorem 1.8. Then observe that, after dividing by p , the derivative of Φ_j^{\log} induces a morphism

$$\delta_j : \Phi_j^* \Omega_{S^{\log}} \rightarrow \Omega_{S^{\log}}$$

We claim that δ_j is an *isomorphism*. Indeed, this follows immediately from Lemma 2.3 (cf. also [Mzk1], Chapter III, Propositions 2.1, 2.3). We state this as a Theorem:

Theorem 3.2. *Let Π be a binary VF-pattern of period ϖ . Let S^{\log} and $\Phi_j^{\log} : S^{\log} \rightarrow S^{\log}$ be as in Theorem 1.8. Then the morphism*

$$\delta_j : \Phi_j^* \Omega_{S^{\log}} \rightarrow \Omega_{S^{\log}}$$

given by dividing the derivative of Φ_j^{\log} by p is an isomorphism. We shall call refer to this property of the system of Frobenius liftings $\{\Phi_j^{\log}\}$ by saying that this system of Frobenius liftings is Π -ordinary.

§3.2. The Spiked Case

Because of the technical difficulties involved, in the present work, we will only consider the case $\varpi = 2$ in detail. Thus, let Π be a spiked VF-pattern of period $\varpi = 2$. Let us write c for the unique nonnegative integer such that $\Pi(1) = \chi - \frac{1}{2} \cdot c$. We shall call c the *colevel* of Π . Next, let us assume that we are given an étale morphism

$$S \rightarrow (\overline{\mathcal{N}}_{g,r}^{\Pi,s})^{\text{ord}}$$

Thus, S is a smooth \mathbf{F}_p -scheme. Equip S with the log structure pulled back from $\overline{\mathcal{M}}_{g,r}^{\log}$. Let

$$f^{\log} : X^{\log} \rightarrow S^{\log}$$

be the tautological log-curve pulled back from $\overline{\mathcal{M}}_{g,r}^{\log}$. Let

$$\mathbf{P}_0, \mathbf{P}_1$$

be the tautological 0th and 1st link bundles on X^{\log} and $(X^{\log})^F$ (pulled back from $\overline{\mathcal{N}}_{g,r}^{\Pi,s}$), respectively. Thus, \mathbf{P}_0 is an indigenous bundle on X^{\log} , and \mathbf{P}_1 is crys-stable of level $\Pi(1) > 0$ on $(X^{\log})^F$. Let

$$D \subseteq X^F$$

be the *p-curvature locus* (cf. Chapter III, Definition 1.2, (3)) of the indigenous bundle \mathbf{P}_0 . Thus, D is a finite, flat S -scheme of constant degree c .

Let us review what it means for S to map into the Π -ordinary locus (cf. Definition 1.2). Write

$$\mathcal{S} \stackrel{\text{def}}{=} (S^{\log})^{\text{pro}}$$

(cf. Chapter VI, Definition 1.9), i.e., the “perfection” of S^{\log} . Then in the discussion preceding Definition 1.1, we considered various vector bundles “ $\Theta(\mathcal{P}_i)$ ” and “ $\Theta'(\mathcal{P}_i)$ ” (for $i = 0, 1$). Let us denote the bundle on \mathcal{S} corresponding to “ $\Theta(\mathcal{P}_i)$ ” (respectively, “ $\Theta'(\mathcal{P}_i)$ ”) by $(\Theta_i)_\mathcal{S}$ (respectively, $(\Theta'_i)_\mathcal{S}$). Recall, moreover, that we had various exact sequences

$$0 \rightarrow (F^0/F^1)_\mathcal{S} \rightarrow (\Theta_1)_\mathcal{S} \rightarrow (F^{-1}/F^0)_\mathcal{S} \rightarrow 0$$

and

$$0 \rightarrow (F^{-1}/F^0)_\mathcal{S} \rightarrow (\Theta'_1)_\mathcal{S} \rightarrow (F^0/F^1)_\mathcal{S} \rightarrow 0$$

Moreover, the assumption of Π -ordinariness implied that the Frobenius action on the de Rham cohomology modules of \mathbf{P}_0 and \mathbf{P}_1 gave rise to *isomorphisms*:

$$(\Phi_0)_\mathcal{S} : \Phi_\mathcal{S}^*(\Theta_0)_\mathcal{S} \cong (\Theta_1)_\mathcal{S}$$

and

$$(\Phi_1)_\mathcal{S} : \Phi_\mathcal{S}^*(\Theta'_1)_\mathcal{S} \cong (\Theta_0)_\mathcal{S}$$

Next, we would like to consider the *p-curvature locus* $D \subseteq X^F$ in greater detail. First, let us recall the moduli spaces “ $\mathcal{D}_{X/S}^l$ ” of “ l -balanced divisors” discussed in Chapter I, Definition 3.7. Let us write \mathcal{D} and \mathcal{D}^F for the spaces (i.e., the “ $\mathcal{D}_{X/S}^l$ ”) of $\Pi(1)$ -balanced divisors on

X^{\log} and $(X^{\log})^F$, respectively. Since the p -curvature locus is necessarily $\Pi(1)$ -balanced (cf. Chapter III, Definition 1.2, (3)), it defines a section

$$\sigma : S \rightarrow \mathcal{D}^F$$

of $\mathcal{D}^F \rightarrow S$. Since \mathcal{D} is smooth over S (of relative dimension c), and \mathcal{D}^F is the result of base-changing \mathcal{D} by the Frobenius morphism $\Phi_S : S \rightarrow S$, it follows that the tangent bundle $\Theta_{\mathcal{D}^F}$ of \mathcal{D}^F splits naturally as a direct sum:

$$\Theta_{\mathcal{D}^F} = \Theta_S \oplus \Theta_{\mathcal{D}^F/S}$$

Thus, differentiating σ gives a morphism

$$\Theta_\sigma : \Theta_S \rightarrow \sigma^* \Theta_{\mathcal{D}^F/S}$$

Note that Θ_σ is *surjective*. Indeed, if L_0 is the $(\chi, \Pi(1))$ -link stack $\mathcal{R}_{g,r}^{\chi, \Pi(1)}$ (cf. Chapter III, Definition 1.2), then the étale morphism $S \rightarrow (\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$ factors through L_0 (indeed, consider the classifying morphism for the bundle $\mathbf{P}_0!$). Moreover, since L_0 is *smooth* over \mathbf{F}_p of dimension $3g - 3 + r$ (cf. Chapter III, Corollary 1.6), it follows that the resulting morphism $S \rightarrow L_0$ is also *étale*. Now observe that (from the definition of L_0) Θ_σ factors through $\Theta_{L_0|S}$ and that the resulting morphism $\Theta_{L_0|S} \rightarrow \sigma^* \Theta_{\mathcal{D}^F/S}$ is *surjective* (by the proof of Chapter III, Corollary 1.6 – see, especially, the paragraph concerning the case “ $j > 0$ ”). This shows that Θ_σ is surjective, as desired. In fact, this same proof even shows that the morphism

$$\Theta_\sigma^{\log} : \Theta_{S^{\log}} \rightarrow \sigma^* \Theta_{\mathcal{D}^F/S}$$

obtained by restricting Θ_σ to $\Theta_{S^{\log}} \subseteq \Theta_S$ is *surjective*, as well.

Next, let T be the scheme defined by the cartesian diagram:

$$\begin{array}{ccc} T & \longrightarrow & S \\ \downarrow & & \downarrow \sigma \\ \mathcal{D} & \xrightarrow{\Phi_{\mathcal{D}/S}} & \mathcal{D}^F \end{array}$$

and let

$$E \subseteq X_T \stackrel{\text{def}}{=} X \times_S T$$

be the divisor defined by the morphism $T \rightarrow \mathcal{D}$. Thus, the morphism $E \rightarrow T$ is finite and flat of degree c , and maps by the relative Frobenius morphism $\Phi_{X_T/T}$ to the divisor $D_T = D \times_S T \subseteq X_T$.

Now let us observe that T is \mathbf{F}_p -smooth. Indeed, the fact that Θ_σ is surjective implies that we may choose local coordinates s_1, \dots, s_d (where $d \stackrel{\text{def}}{=} 3g - 3 + r \geq c$) on S and relative local coordinates $\delta_1, \dots, \delta_c$ of \mathcal{D}^F over S such that the image of σ is defined by the equations $s_1 - \delta_1, \dots, s_c - \delta_c$. Moreover, \mathcal{D} is defined locally “over” (i.e., relative to $\Phi_{\mathcal{D}/S}$) \mathcal{D}^F by adjoining τ_1, \dots, τ_c satisfying $\tau_1^p - \delta_1, \dots, \tau_c^p - \delta_c$ to $\mathcal{O}_{\mathcal{D}^F}$. Thus, T is defined over S by adjoining t_1, \dots, t_c (i.e., the images of τ_1, \dots, τ_c) satisfying $t_1^p - s_1, \dots, t_c^p - s_c$ to \mathcal{O}_S . That is to say, $t_1, \dots, t_c, s_{c+1}, \dots, s_d$ forms a local system of coordinates on T . This completes the proof that T is \mathbf{F}_p -smooth. In fact, the surjectivity of Θ_σ^{\log} even implies that T^{\log} (with the log structure pulled back from S) is log smooth over \mathbf{F}_p .

Next, let us observe that it follows from the definition of T as a fibered product (or from the above explicit analysis in terms of local coordinates) that the Frobenius morphism

$$\Phi_{S^{\log}} : S^{\log} \rightarrow S^{\log}$$

factors through T^{\log} . Thus, we obtain morphisms

$$\Phi^{\text{Hs}, \log} : S^{\log} \rightarrow T^{\log}; \quad \Phi^{\text{spk}, \log} : T^{\log} \rightarrow S^{\log}$$

such that $\Phi_{S^{\log}} = \Phi^{\text{spk}, \log} \circ \Phi^{\text{Hs}, \log}$. Here, “Hs” (respectively, “spk”) stands for “Hodge section” (respectively, “spikes”) since this portion of the Frobenius morphism corresponds to taking the p^{th} roots of coordinates that parametrize the obstruction to deforming the Hodge section of \mathbf{P}_1 (respectively, parametrize the spikes, i.e., deformations of D). Then taking derivatives gives us morphisms of tangent bundles

$$\delta^{\text{spk}} : \Theta_{T^{\log}} \rightarrow (\Phi^{\text{spk}})^* \Theta_{S^{\log}}$$

and

$$\delta^{\text{Hs}} : \Theta_{S^{\log}} \rightarrow (\Phi^{\text{Hs}})^* \Theta_{T^{\log}}$$

Moreover, the explicit local analysis involving coordinates in the preceding paragraph implies that Φ^{spk} (respectively, Φ^{Hs}) essentially corresponds to taking p^{th} roots of the coordinates s_1, \dots, s_c (respectively, s_{c+1}, \dots, s_d). Thus, we see that we have exact sequences

$$\Theta_{S^{\log}} \xrightarrow{\delta^{\text{Hs}}} (\Phi^{\text{Hs}})^* \Theta_{T^{\log}} \xrightarrow{(\Phi^{\text{Hs}})^* \delta^{\text{spk}}} \Phi_S^* \Theta_{S^{\log}}$$

and

$$\Theta_{T^{\log}} \xrightarrow{\delta^{\text{spk}}} (\Phi^{\text{spk}})^* \Theta_{S^{\log}} \xrightarrow{(\Phi^{\text{spk}})^* \delta^{\text{Hs}}} \Phi_T^* \Theta_{T^{\log}}$$

We summarize this as follows:

Lemma 3.3. *Let T be as in the cartesian diagram above. Equip T with the log structure obtained by pulling back the log structure of S . Then T (respectively, T^{\log}) is smooth (respectively, log smooth) over \mathbf{F}_p . Moreover, the Frobenius morphism $\Phi_{S^{\log}} : S^{\log} \rightarrow S^{\log}$ factors as $\Phi^{\text{spk}, \log} \circ \Phi^{\text{Hs}, \log}$, where*

$$\Phi^{\text{Hs}, \log} : S^{\log} \rightarrow T^{\log}, \quad \Phi^{\text{spk}, \log} : T^{\log} \rightarrow S^{\log}$$

are morphisms whose derivatives induce exact sequences

$$\Theta_{S^{\log}} \xrightarrow{\delta^{\text{Hs}}} (\Phi^{\text{Hs}})^* \Theta_{T^{\log}} \xrightarrow{(\Phi^{\text{Hs}})^* \delta^{\text{spk}}} \Phi_S^* \Theta_{S^{\log}}$$

and

$$\Theta_{T^{\log}} \xrightarrow{\delta^{\text{spk}}} (\Phi^{\text{spk}})^* \Theta_{S^{\log}} \xrightarrow{(\Phi^{\text{spk}})^* \delta^{\text{Hs}}} \Phi_T^* \Theta_{T^{\log}}$$

of vector bundles on S .

Next, we would like to consider the bundle $\mathbf{P}_1|_T$ on $(X_T^{\log})^F$. (In the following discussion, we use the notation “ $|_T$ ” to denote pull-back by Φ^{spk} .) Write L_1 for the $(\Pi(1), \chi)$ -link stack $\mathcal{R}_{g,r}^{\Pi(1), \chi}$ (cf. Chapter III, Definition 1.2). Note that $\mathbf{P}_1|_T$ defines a classifying morphism

$$\kappa^F : T \rightarrow L_1$$

Moreover, the pull-back to T of the log tangent bundle of L_1 (equipped with the log structure pulled back from $\overline{\mathcal{M}}_{g,r}^{\log}$) via this classifying morphism may be identified with

$$(\Phi_S^* \Theta_{S^{\log}})^F|_T = (\Phi_S^*)^2 \Theta_{S^{\log}}|_T$$

(cf. the proof of Chapter III, Corollary 1.6 in the case “ $j > 0$ ” – i.e., in the notation of *loc. cit.*, the tangent bundle in question parametrizes infinitesimal deformations of the section “ a_0 ” to a section “ a ” of the torsor “ \mathcal{A} ”). Thus, taking the derivative of the classifying morphism gives us a morphism

$$\Theta_{\kappa^F} : \Theta_{T^{\log}} \rightarrow (\Phi_S^*)^2 \Theta_{S^{\log}}|_T$$

Now I *claim* that:

If $\Theta_{\kappa^F} = 0$, then $\mathbf{P}_1|_T$ is the pull-back via Φ_T of some bundle \mathbf{Q}_1 on X_T^{\log} .

Indeed, to say that $\Theta_{\kappa^F} = 0$ implies that the exterior derivative “d” of the pull-back of any function on L_1 to T via κ^F is zero. But this implies that the pulled back function is the p^{th} power of a function on T , which is another way of saying that the classifying morphism κ^F factors through the Frobenius morphism on T . This completes the proof of the claim.

Let us prove that $\Theta_{\kappa^F} = 0$. To do this, let us first note that the bundle $(\Theta_1)_S$ discussed at the beginning of this subsection descends naturally to some bundle $(\Theta_1)_T^F$. Put another way, $(\Theta_1)_T^F$ is the subquotient “ F^{-1}/F^1 ” of the first de Rham cohomology module of the bundle $\mathbf{P}_1|_T$. Moreover, it follows from the definition of the isomorphism $(\Phi_0)_S$ (cf. the discussion preceding Definition 1.1) that the isomorphism $(\Phi_0)_S$ also admits a “ $\mathbf{P}_1|_T$ -version,” i.e., descends to T . Thus, (since $(\Theta_0)_S$ descends to T in the present context as $\Phi_S^* \Theta_{S^{\log}}|_T$) we get an isomorphism

$$(\Phi_0)_T^F : (\Phi_S^*)^2 \Theta_{S^{\log}}|_T \rightarrow (\Theta_1)_T^F$$

which we may compose with Θ_{κ^F} to obtain a morphism

$$\Theta_{\mathbf{P}_1|_T} : \Theta_{T^{\log}} \rightarrow (\Theta_1)_T^F$$

Since $(\Phi_0)_T^F$ is an *isomorphism*, in order to show that $\Theta_{\kappa^F} = 0$, it suffices to show that $\Theta_{\mathbf{P}_1|_T} = 0$.

On the other hand, we may interpret $\Theta_{\mathbf{P}_1|_T}$ *functorially* as follows: To do this, we begin by observing that $(\Theta_1)_T^F$ fits into an exact sequence

$$0 \rightarrow F^0/F^1((\Theta_1)_T^F) \rightarrow (\Theta_1)_T^F \rightarrow F^{-1}/F^0((\Theta_1)_T^F) \rightarrow 0$$

Let us denote by

$$\Theta_{\mathbf{P}_1|_T}^{-1} : \Theta_{T^{\log}} \rightarrow F^{-1}/F^0((\Theta_1)_T^F)$$

the composite of $\Theta_{\mathbf{P}_1|_T}$ with the projection to $F^{-1}/F^0((\Theta_1)_T^F)$. But it follows from Chapter I, Proposition 1.7 (cf. also the proof of Chapter

I, Lemma 3.8) that $\Theta_{\mathbf{P}_1|T}^{-1}$ may be interpreted in the following way: First of all, $F^{-1}/F^0((\Theta_1)_T^F)$ is naturally a quotient of

$$\Theta_{S^{\log}|T}^F = (\Phi_S^* \Theta_{S^{\log}})|_T = (\Phi_S^* \Theta_{\overline{\mathcal{M}}_{g,r}^{\log}|S})|_T$$

which may be thought of as the bundle that carries the obstruction to deforming the Hodge section of the crys-stable bundle $\mathbf{P}_1|_T$. Indeed, if $f^{\log} : X_T^{\log} \rightarrow T^{\log}$ is the structure morphism for the curve X_T^{\log} , then this quotient may be identified with the pull-back by Φ_T of the natural quotient

$$\mathbf{R}^1 f_{\mathrm{DR},*}(\tau_{X_T^{\log}/T^{\log}}) \rightarrow \mathbf{R}^1 f_{\mathrm{DR},*}(\tau_{X_T^{\log}/T^{\log}}(E))$$

(i.e., the bundle “ \mathcal{L}^{-1} ” of Chapter I, Proposition 1.7, is $\tau_{X_T^{\log}/T^{\log}}(E)$). Thus, by thinking of $\Theta_{\mathbf{P}_1|T}^{-1}$ as the morphism that associates to an infinitesimal deformation on T the corresponding obstruction to deforming the Hodge section of $\mathbf{P}_1|_T$, it follows that $\Theta_{\mathbf{P}_1|T}^{-1}$ is none other than the composite of the Kodaira-Spencer morphism

$$\Theta_{T^{\log}} \rightarrow \Theta_{S^{\log}|T}^F$$

for the curve $(X_T^{\log})^F$ with the surjection $\Theta_{S^{\log}|T}^F \rightarrow F^{-1}/F^0((\Theta_1)_T^F)$. On the other hand, this Kodaira-Spencer morphism is *zero* (since the classifying morphism $T \rightarrow \overline{\mathcal{M}}_{g,r}$ for the curve $(X_T^{\log})^F$ factors through the Frobenius on T). This shows that $\Theta_{\mathbf{P}_1|T}^{-1} = 0$.

Thus, we see that the morphism $\Theta_{\mathbf{P}_1|T} : \Theta_{T^{\log}} \rightarrow (\Theta_1)_T^F$ defines a morphism

$$\Theta_{\mathbf{P}_1|T}^0 : \Theta_{T^{\log}} \rightarrow F^0/F^1((\Theta_1)_T^F)$$

But now, by Chapter I, Lemma 3.8, it follows that this morphism may be interpreted as the morphism that associates to an infinitesimal deformation on $\Theta_{T^{\log}}$ the corresponding deformation in the “Kodaira-Spencer locus” (cf. Chapter I, Definition 3.6) $E^F \subseteq X_T^F$ (i.e., the divisor where the Kodaira-Spencer morphism associated to the Hodge section of $\mathbf{P}_1|_T$ is zero). On the other hand, since $E^F \subseteq X_T^F$ is the pull-back by Φ_T of a divisor $E \subseteq X_T$, it follows that this derivative is *zero*, i.e., $\Theta_{\mathbf{P}_1|T}^0 = 0$. This completes the proof that $\Theta_{\mathbf{P}_1|T} = 0$, hence also the proof that $\Theta_{\kappa^F} = 0$. In particular, we see that we have proven the following result:

Lemma 3.4. *The crys-stable bundle \mathbf{P}_1 on $(X^{\log})^F$ is the pull-back via Φ^{Hs} of some crys-stable bundle \mathbf{Q}_1 (also of level $\Pi(1)$) on X_T^{\log} .*

Proof. By the “claim” proven above (in the paragraph following Lemma 3.3), we see that we have shown that $\mathbf{P}_1|_T$ descends to some \mathbf{Q}_1 on X_T^{\log} . Thus, $(\Phi^{\text{Hs}})^*(\mathbf{Q}_1)$ and \mathbf{P}_1 become isomorphic after pull-back to T^{\log} . Since, however, $\Phi^{\text{spk}} : T \rightarrow S$ is *faithfully flat*, and the space of crystal stable bundles admits a *fine moduli stack in the category of algebraic spaces* (Chapter I, Theorem 2.7), it follows that $(\Phi^{\text{Hs}})^*$ and \mathbf{P}_1 are already isomorphic on $(X^{\log})^F$, as desired. \circ

Next, we would like to consider the various derivatives of classifying morphisms for \mathbf{Q}_1 that we considered above for $\mathbf{P}_1|_T$. First, we observe that \mathbf{Q}_1 defines a classifying morphism

$$\kappa : T \rightarrow L_1$$

to the link stack L_1 . Its derivative is a morphism

$$\Theta_\kappa : \Theta_{T^{\log}} \rightarrow (\Phi_S^* \Theta_{S^{\log}})|_T$$

Let us write Θ_T^{spk} (respectively, Θ_T^{Hs}) for the kernel (respectively, image) of the morphism of vector bundles $\delta^{\text{spk}} : \Theta_{T^{\log}} \rightarrow \Theta_{S^{\log}}|_T$. Thus, we have an exact sequence

$$0 \rightarrow \Theta_T^{\text{spk}} \rightarrow \Theta_{T^{\log}} \rightarrow \Theta_T^{\text{Hs}} \rightarrow 0$$

Moreover, the isomorphism $(\Phi_0)_T^F$ considered above clearly descends to an isomorphism

$$(\Phi_0)_T : (\Phi_S^* \Theta_{S^{\log}})|_T \rightarrow (\Theta_1)_T$$

on T . Next, observe that $(\Theta_1)_T$ fits into an exact sequence

$$0 \rightarrow F^0/F^1((\Theta_1)_T) \rightarrow (\Theta_1)_T \rightarrow F^{-1}/F^0((\Theta_1)_T) \rightarrow 0$$

which descends the analogous exact sequence for $(\Theta_1)_T^F$ considered above. As reviewed above (cf. Chapter I, Lemma 3.8), the subbundle $F^0/F^1((\Theta_1)_T) \subseteq (\Theta_1)_T$ corresponds to deformations of the Kodaira-Spencer locus of \mathbf{Q}_1 (i.e., of $E \subseteq X_T$). Thus, it follows from our explicit analysis of T involving local coordinates that the morphism

$$\Theta_{\mathbf{Q}_1} : \Theta_{T^{\log}} \rightarrow (\Theta_1)_T$$

obtained by composing Θ_κ with the isomorphism $(\Phi_0)_T$ maps the rank c subbundle $\Theta_T^{\text{spk}} \subseteq \Theta_{T^{\log}}$ isomorphically onto the rank c subbundle

$F^0/F^1((\Theta_1)_T) \subseteq (\Theta_1)_T$. In other words, we have a morphism of exact sequences:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \Theta_T^{\text{spk}} & \longrightarrow & \Theta_{T^{\log}} & \longrightarrow & \Theta_T^{\text{Hs}} & \longrightarrow & 0 \\ & & \downarrow \alpha & & \downarrow \Theta_{\mathbf{Q}_1} & & \downarrow \beta & & \\ 0 & \longrightarrow & F^0/F^1((\Theta_1)_T) & \longrightarrow & (\Theta_1)_T & \longrightarrow & F^{-1}/F^0((\Theta_1)_T) & \longrightarrow & 0 \end{array}$$

where α and β are defined so as to make the diagram commutative, and α is an isomorphism. In particular, we may read off the following result from this morphism of exact sequences:

Lemma 3.5. *The following four conditions are equivalent:*

(i.) $\Theta_\kappa : \Theta_{T^{\log}} \rightarrow (\Phi_S^* \Theta_{S^{\log}})|_T$ is an isomorphism.

(ii.) $\Theta_{\mathbf{Q}_1} : \Theta_{T^{\log}} \rightarrow (\Theta_1)_T$ is an isomorphism.

(iii.) $\beta : \Theta_T^{\text{Hs}} \rightarrow F^{-1}/F^0((\Theta_1)_T)$ is an isomorphism.

(iv.) The morphism $\gamma : \Theta_{T^{\log}} \rightarrow F^{-1}/F^0((\Theta_1)_T)$ (obtained from the above morphism of exact sequences) is surjective.

We are now ready to make the following important definition:

Definition 3.6. We will refer to the open substack of $(\overline{\mathcal{N}}_{g,r}^{\Pi,s})^{\text{ord}}$ on which:

(a.) The four equivalent conditions of Lemma 3.5 are satisfied.

(b.) The composite morphism

$$\Phi_T^*(\Theta_T^{\text{Hs}}) \longrightarrow \Phi_T^*(\Theta_{S^{\log}}|_T) \xrightarrow{(\Phi_0)_T} (\Theta_1)_T \longrightarrow F^{-1}/F^0((\Theta_1)_T)$$

(where the first and last morphisms are the natural inclusion and projection, respectively) is an isomorphism.

as the *very (II-) ordinary locus* of $\overline{\mathcal{N}}_{g,r}^{\Pi,s}$.

We conclude this subsection with existence results. Let $p \geq 5$. First, we consider the case $(g, r) = (0, 5)$. Then the data constructed

in the discussion preceding Chapter IV, Lemma 3.6, *trivially* satisfies conditions (a.) and (b.) of Definition 3.6 above. Indeed, this follows from the fact that both Θ_T^{Hs} and $F^{-1}/F^0((\Theta_1)_T)$ are vector bundles of rank $(3g - 3 + r) - c = 2 - 2 = 0$ (since $c = 2$) on T .

Theorem 3.7. *Suppose that $2g - 2 + r \geq 3$ and $p \geq 5$. Then there exists a spiked VF-pattern Π of period 2 such that the very Π -ordinary locus of $\overline{\mathcal{N}}_{g,r}^{\Pi,s}$ is nonempty.*

Proof. The case of arbitrary (g, r) (such that $2g - 2 + r \geq 3$) is derived from the $(g, r) = (0, 5)$ case, just as Chapter IV, Theorem 3.7, is derived from Chapter IV, Lemma 3.6, i.e., via gluing. The nodes at which the gluing takes place are, so to speak, “classical ordinary” (i.e., they have radius equal to 0 – cf. the usage of this term in Chapter V); moreover, it is the new modular coordinates corresponding to these nodes that give rise to the “Hs” coordinates of the discussion of the present subsection. Thus, the Frobenius action (hence, in particular, the morphism β of condition (iii) of Lemma 3.5 and the composite morphism of condition (b.) of Definition 3.6) on these coordinates will be an *isomorphism* (as is always the case for “classical ordinary nodes” – cf., e.g., the calculation in the discussion preceding [Mzk1], Chapter II, Proposition 3.7). In particular, conditions (a.) and (b.) of Definition 3.6 are satisfied. This completes the proof. \circ

§3.3. **Frobenius Liftings in the Very Ordinary Case**

We maintain the notation of §3.2. Assume, moreover, that

$$S \rightarrow (\overline{\mathcal{N}}_{g,r}^{\Pi,s})^{\text{very ord}}$$

is an étale morphism into the *very Π -ordinary locus* (cf. Definition 3.6). Thus, S is étale over $\overline{\mathcal{M}}_{g,r}$; equip S with the log structure pulled back from that of $\overline{\mathcal{M}}_{g,r}^{\text{log}}$, and call the resulting log stack S^{log} . By Theorem 3.7 above, we know that if p , g , and r satisfy certain conditions, then such a (nonempty) S exists.

Let

$$\Phi^{\text{Hs},\text{log}} : S^{\text{log}} \rightarrow T^{\text{log}}; \quad \Phi^{\text{spk},\text{log}} : T^{\text{log}} \rightarrow S^{\text{log}}$$

be as in Lemma 3.3. Since S is étale over $(\overline{\mathcal{M}}_{g,r})_{\mathbf{F}_p}$, it admits a natural lifting $S_{\mathbf{Z}_p}$ which is étale over $(\overline{\mathcal{M}}_{g,r})_{\mathbf{Z}_p}$. Since T^{log} is log smooth over \mathbf{F}_p , all liftings of T^{log} to \mathbf{Z}_p are étale locally (on T) isomorphic; suppose that we are given one fixed such lifting $T_{\mathbf{Z}_p}^{\text{log}}$.

In Theorem 1.8 (under the assumption of Π -ordinariness), we constructed a certain *canonical Frobenius lifting*:

$$\Phi^{\log} : S_{\mathbf{Z}_p}^{\log} \rightarrow S_{\mathbf{Z}_p}^{\log}$$

whose reduction modulo p is the square of the absolute Frobenius on S . The purpose of this subsection is to observe that:

Under the assumption of very ordinariness (in fact, condition (a.) of Definition 3.6 is sufficient), this Frobenius lifting Φ^{\log} admits a natural decomposition into a composite of two Frobenius liftings.

The proof of the existence of these two Frobenius liftings is similar to the proof of Theorem 1.8. The main difference, however, is that whereas $S_{\mathbf{Z}_p}$ is étale over $(\overline{\mathcal{M}}_{g,r})_{\mathbf{Z}_p}$, which has a *modular interpretation* (i.e., as the moduli stack of certain types of curves), $T_{\mathbf{Z}_p}$ does not admit an étale morphism to any sort of natural moduli space. This causes problems because, whereas in the case of Theorem 1.8, we constructed Frobenius liftings on $S_{\mathbf{Z}_p}$ by using the modular interpretation of (deformations of morphisms to) $S_{\mathbf{Z}_p}$, here, we wish to construct certain morphisms to $T_{\mathbf{Z}_p}$ (via deformation theory) despite the fact that we do not have a modular interpretation for (deformations of morphisms to) $T_{\mathbf{Z}_p}$.

Let us write

$$\theta^{\log} \stackrel{\text{def}}{=} \Phi^{\text{Hs}, \log} : S^{\log} \rightarrow T^{\log}; \quad \phi^{\log} \stackrel{\text{def}}{=} \Phi_S^{\log} \circ \Phi^{\text{spk}, \log} : T^{\log} \rightarrow S^{\log}$$

(where Φ_S^{\log} is the usual Frobenius morphism on the \mathbf{F}_p -log scheme S^{\log}). Thus, ϕ^{\log} is, so to speak, the square of Frobenius (i.e., the “ p^2 -power map”) with respect to the “spk” coordinates, and the usual Frobenius morphism (i.e., the “ p -power map”) with respect to the “Hs” coordinates, while θ^{\log} is Frobenius with respect to the “Hs” coordinates, and an isomorphism with respect to the “spk” coordinates. In particular, we have $\phi^{\log} \circ \theta^{\log} = (\Phi_S^{\log})^2$.

Next, observe that the “first link bundle” \mathbf{Q}_1 on X_T^{\log} is nilpotent and admissible – i.e., defined by a section of the torsor “ \mathcal{A} ” of Chapter III, Proposition 1.5, in the case where the divisor “ D ” of *loc. cit.* is the empty set. Moreover, by the relationship between Frobenius liftings and “FL-bundles” (cf. [Mzk1], Chapter II, §1, especially Proposition 1.2), such sections of the torsor “ \mathcal{A} ” are in natural bijective correspondence with liftings of ϕ^{\log} to $\mathbf{Z}/p^2\mathbf{Z}$. Thus, we obtain that \mathbf{Q}_1 defines a natural lifting

$$\phi_{\mathbf{Z}/p^2\mathbf{Z}}^{\log} : T_{\mathbf{Z}/p^2\mathbf{Z}}^{\log} \rightarrow S_{\mathbf{Z}/p^2\mathbf{Z}}^{\log}$$

of ϕ^{\log} to $\mathbf{Z}/p^2\mathbf{Z}$. Another way to put this is that $\phi_{\mathbf{Z}/p^2\mathbf{Z}}^{\log}$ is the unique lifting with the following property:

Given any lifting $(\mathbf{P}_0)_{\mathbf{Z}/p^2\mathbf{Z}}$ (of \mathbf{P}_0 to an indigenous bundle on the tautological curve $X_{\mathbf{Z}/p^2\mathbf{Z}}^{\log} \rightarrow S_{\mathbf{Z}/p^2\mathbf{Z}}^{\log}$), if we use the “Frobenius lifting” $\phi_{\mathbf{Z}/p^2\mathbf{Z}}^{\log}$ to form the renormalized Frobenius pull-back $\mathbf{F}^*((\mathbf{P}_0)_{\mathbf{Z}/p^2\mathbf{Z}})_{\mathbf{F}_p}$ – which will be a bundle on X_T^{\log} – then this bundle $\mathbf{F}^*((\mathbf{P}_0)_{\mathbf{Z}/p^2\mathbf{Z}})_{\mathbf{F}_p}$ is equal to \mathbf{Q}_1 .

Thus, at this point, the proof looks just like the proof of Theorem 1.8.

Next, we would like to lift $\theta^{\log} : S^{\log} \rightarrow T^{\log}$ to $\mathbf{Z}/p^2\mathbf{Z}$. In order to construct such a lifting, we must find a modular interpretation for deformations of morphisms to $T_{\mathbf{Z}_p}^{\log}$. The interpretation that we will use is the following:

We will think of $T_{\mathbf{Z}_p}^{\log}$ as a certain $(3g - 3 + r)$ -dimensional subspace of the $(2(3g - 3 + r))$ -dimensional space of deformations of the crystal (on $\text{Crys}(X_T^{\log}/T^{\log})$ defined by) \mathbf{Q}_1 to a crystal on $\text{Crys}(X_T^{\log}/T_{\mathbf{Z}_p}^{\log})$.

More concretely, we propose to do the following: Suppose that (for some $n \geq 0$) we have chosen a crystal $(\mathbf{Q}_1)_{\mathbf{Z}/p^{n+1}\mathbf{Z}}$ on $\text{Crys}(X_T^{\log}/(T_{\mathbf{Z}/p^{n+1}\mathbf{Z}})^{\log})$ lifting \mathbf{Q}_1 . Note that this “rigidifies” $T_{\mathbf{Z}_p}$ modulo p^{n+1} in the sense that there do not exist any automorphisms of $T_{\mathbf{Z}_p}$ which stabilize $(\mathbf{Q}_1)_{\mathbf{Z}/p^{n+1}\mathbf{Z}}$ and are the identity modulo p (cf. the fact that the morphism $\Theta_{\mathbf{Q}_1}$ of Lemma 3.5, (ii.), is an *isomorphism*). Then if choose a crystal $(\mathbf{Q}_1)_{\mathbf{Z}/p^{n+2}\mathbf{Z}}$ on $\text{Crys}(X_T^{\log}/(T_{\mathbf{Z}/p^{n+2}\mathbf{Z}})^{\log})$ lifting $(\mathbf{Q}_1)_{\mathbf{Z}/p^{n+1}\mathbf{Z}}$, we get a rigidification of $T_{\mathbf{Z}_p}$ modulo p^{n+2} . In fact, the fact that the morphism $\Theta_{\mathbf{Q}_1}$ of Lemma 3.5, (ii.), is an *isomorphism* implies that this rigidification is determined as soon as $(\mathbf{Q}_1)_{\mathbf{Z}/p^{n+2}\mathbf{Z}}$ is determined modulo $p^{n+1} \cdot F^1$ (cf. the terminology of §1.3). (Note: the vector bundle $(\Theta_1)_T$ of Lemma 3.5, (ii.), may be identified with the Hodge filtration quotient F^{-1}/F^1 of the first de Rham cohomology module of \mathbf{Q}_1). Finally, we remark that such liftings $(\mathbf{Q}_1)_{\mathbf{Z}/p^{n+2}\mathbf{Z}}$ always exist étale locally on $S_{\mathbf{Z}_p}$ or $T_{\mathbf{Z}_p}$, and this will be sufficient for us since our construction will be canonical, thus allowing us to glue together objects constructed locally.

Next, observe that (by elementary deformation theory) *up to automorphisms of $S_{\mathbf{Z}/p^2\mathbf{Z}}^{\log}$ which are the identity modulo p , the liftings of $\theta^{\log} = \Phi^{\text{Hs}, \log} : S^{\log} \rightarrow T^{\log}$ to $\mathbf{Z}/p^2\mathbf{Z}$ form a torsor over the cokernel $\text{Coker}(\delta^{\text{Hs}})$ of $\delta^{\text{Hs}} : \Theta_{S^{\log}} \rightarrow (\Phi^{\text{Hs}})^*\Theta_{T^{\log}}$. Moreover, it follows from the first exact sequence of Lemma 3.3 that this cokernel $\text{Coker}(\delta^{\text{Hs}})$ may be identified with $(\Phi^{\text{Hs}})^*\Theta_T^{\text{Hs}}$. Thus, if we compose this identification with $(\Phi^{\text{Hs}})^*(\beta)$ (where β is the isomorphism of Lemma 3.5), we obtain an isomorphism*

$$\zeta : \text{Coker}(\delta^{\text{Hs}} : \Theta_{S^{\log}} \rightarrow (\Phi^{\text{Hs}})^* \Theta_{T^{\log}}) \cong (\Phi^{\text{Hs}})^*(F^{-1}/F^0((\Theta_1)_T))$$

On the other hand, the zeroth link bundle \mathbf{P}_0 on S arises from a section of the torsor “ \mathcal{A} ” of Chapter III, Proposition 1.5, in the case where the “ D ” of *loc. cit.* is equal to $D \subseteq X^F$ of the present discussion. Moreover, the vector bundle $(\Phi^{\text{Hs}})^*(F^{-1}/F^0((\Theta_1)_T))$ is easily seen to be the same as the vector bundle “ Θ_D ” of Chapter III, Proposition 1.5. Thus, the fact that ζ is an isomorphism implies that there exists a unique (i.e., up to automorphisms of $S_{\mathbf{Z}/p^2\mathbf{Z}}^{\log}$ which are the identity modulo p) lifting

$$\theta_{\mathbf{Z}/p^2\mathbf{Z}}^{\log} : S_{\mathbf{Z}/p^2\mathbf{Z}}^{\log} \rightarrow T_{\mathbf{Z}/p^2\mathbf{Z}}^{\log}$$

corresponding to the section of “ \mathcal{A} ” defined by \mathbf{P}_0 . Another way to put this is the following: Suppose that we have chosen a lifting $(\mathbf{Q}_1)_{\mathbf{Z}/p^2\mathbf{Z}}$ of \mathbf{Q}_1 modulo $p \cdot F^1$. Then $\theta_{\mathbf{Z}/p^2\mathbf{Z}}^{\log}$ is the unique lifting (up to composition with an automorphism of $S_{\mathbf{Z}/p^2\mathbf{Z}}^{\log}$ which is the identity modulo p) to $\mathbf{Z}/p^2\mathbf{Z}$ such that the “renormalized Frobenius pull-back” $\mathbf{F}^*((\mathbf{Q}_1)_{\mathbf{Z}/p^2\mathbf{Z}})$ that it defines is equal to \mathbf{P}_0 .

Next, we would like to rigidify the choice of lifting $\theta_{\mathbf{Z}/p^2\mathbf{Z}}^{\log}$ relative to automorphisms of $S_{\mathbf{Z}/p^2\mathbf{Z}}^{\log}$ which are the identity modulo p . Observe that the space of liftings $S_{\mathbf{Z}/p^2\mathbf{Z}}^{\log} \rightarrow T_{\mathbf{Z}/p^2\mathbf{Z}}^{\log}$ which differ by such an automorphism forms a torsor on S over the vector bundle $\text{Im}(\delta^{\text{Hs}})$. Moreover, the first exact sequence of Lemma 3.3 defines an isomorphism between $\text{Im}(\delta^{\text{Hs}})$ and $(\Phi^{\text{Hs}})^* \Theta_T^{\text{spk}}$. Thus, since the morphism α (in the discussion preceding Lemma 3.5) is an isomorphism, we obtain a natural isomorphism

$$\lambda : \text{Im}(\delta^{\text{Hs}}) \cong (\Phi^{\text{Hs}})^*(F^0/F^1((\Theta_1)_T))$$

On the other hand, the quotient $\Theta_{S^{\log}} \rightarrow \text{Im}(\delta^{\text{Hs}})$ may be identified with the quotient $\Theta_{\sigma}^{\log} : \Theta_{S^{\log}} \rightarrow \sigma^* \Theta_{\mathcal{D}^F/S}$ used to construct T (cf. the explicit analysis of T involving local coordinates in §3.2). More intuitively, this quotient of $\Theta_{S^{\log}}$ is the quotient onto the “spk” tangent vectors, i.e., the tangent vectors that correspond to deformations of $D \subseteq X^F$.

Now we would like to define some terminology: We will say that a lifting of \mathbf{P}_0 to a crystal $(\mathbf{P}_0)_{\mathbf{Z}/p^{n+2}\mathbf{Z}}$ on $\text{Crys}(X^{\log}/S_{\mathbf{Z}/p^{n+2}\mathbf{Z}})$ is *spk-indigenous modulo p^{n+1}* if the crystal $(\mathbf{P}_0)_{\mathbf{Z}/p^{n+2}\mathbf{Z}}$ defines an indigenous bundle on the tautological curve $X_{\mathbf{Z}/p^{n+1}\mathbf{Z}}^{\log} \rightarrow S_{\mathbf{Z}/p^{n+1}\mathbf{Z}}^{\log}$, and, moreover, the obstruction to its being indigenous on $X_{\mathbf{Z}/p^{n+2}\mathbf{Z}}^{\log}$ (i.e., the obstruction to lifting the Hodge section modulo p^{n+2}) forms a section of $\Theta_{S^{\log}}$

which vanishes in the quotient $\Theta_{S^{\log}} \rightarrow \text{Im}(\delta^{\text{Hs}})$. Observe, for instance, that *once we fix a lifting* $(\mathbf{Q}_1)_{\mathbf{Z}/p^2\mathbf{Z}}$ *to a crystal on* $\text{Crys}(X_T^{\log}/T_{\mathbf{Z}/p^2\mathbf{Z}})$ *modulo* $p \cdot F^1$, *it makes sense to ask whether or not the renormalized Frobenius pull-back* $\mathbf{F}^*((\mathbf{Q}_1)_{\mathbf{Z}/p^2\mathbf{Z}})$ *defined by a particular* $\theta_{\mathbf{Z}/p^2\mathbf{Z}}^{\log}$ *is spk-indigenous modulo* p . Indeed, this follows from the fact that for different $(\mathbf{Q}_1)_{\mathbf{Z}_p}$'s on $\text{Crys}(X_T^{\log}/T_{\mathbf{Z}_p})$ which coincide modulo $p \cdot F^1$, the resulting $\mathbf{F}^*((\mathbf{Q}_1)_{\mathbf{Z}_p})$'s will coincide modulo p , and their difference modulo p^2 will be in the portion of the first de Rham cohomology module of $\mathbf{F}^*((\mathbf{Q}_1)_{\mathbf{Z}_p})_{\mathbf{F}_p} = \mathbf{P}_0$ which goes to zero under the surjection $\Theta_{S^{\log}} \rightarrow \text{Im}(\delta^{\text{Hs}})$ (where here we regard $\Theta_{S^{\log}} = \Theta_{(\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{F}_p}}|_S$ as the Hodge filtration quotient F^{-1}/F^0 of this de Rham cohomology module).

Thus, by deformation theory, the fact that the morphism λ discussed above is an isomorphism implies that *once we fix a lifting* $(\mathbf{Q}_1)_{\mathbf{Z}/p^2\mathbf{Z}}$ *to a crystal on* $\text{Crys}(X_T^{\log}/T_{\mathbf{Z}/p^2\mathbf{Z}})$ *modulo* $p \cdot F^1$, *there is a unique choice of* $\theta_{\mathbf{Z}/p^2\mathbf{Z}}^{\log}$ *with the property that the corresponding renormalized Frobenius pull-back* $\mathbf{F}^*((\mathbf{Q}_1)_{\mathbf{Z}/p^2\mathbf{Z}})$ *is spk-indigenous modulo* p . We take this as our choice of $\theta_{\mathbf{Z}/p^2\mathbf{Z}}^{\log} : S_{\mathbf{Z}/p^2\mathbf{Z}}^{\log} \rightarrow T_{\mathbf{Z}/p^2\mathbf{Z}}^{\log}$.

Thus, to summarize, we have constructed liftings

$$\theta_{\mathbf{Z}/p^2\mathbf{Z}}^{\log} : S_{\mathbf{Z}/p^2\mathbf{Z}}^{\log} \rightarrow T_{\mathbf{Z}/p^2\mathbf{Z}}^{\log}; \quad \phi_{\mathbf{Z}/p^2\mathbf{Z}}^{\log} : T_{\mathbf{Z}/p^2\mathbf{Z}}^{\log} \rightarrow S_{\mathbf{Z}/p^2\mathbf{Z}}^{\log}$$

for which the corresponding renormalized Frobenius pull-backs of certain bundles possess certain properties. Since $(\mathbf{Q}_1)_{\mathbf{Z}/p^2\mathbf{Z}}$ has been chosen modulo $p \cdot F^1$ and the morphism $(\Theta_0)_S$ (cf. the discussion of §3.2) is an *isomorphism*, it follows (cf. the above discussion, as well as the proof of Theorem 1.8) that there is a unique lifting $\phi_{\mathbf{Z}/p^3\mathbf{Z}}^{\log} : T_{\mathbf{Z}/p^3\mathbf{Z}}^{\log} \rightarrow S_{\mathbf{Z}/p^3\mathbf{Z}}^{\log}$ whose corresponding $\mathbf{F}^*((\mathbf{P}_0)_{\mathbf{Z}/p^3\mathbf{Z}})$ (for some indigenous bundle $(\mathbf{P}_0)_{\mathbf{Z}/p^3\mathbf{Z}}$ on $X_{\mathbf{Z}/p^3\mathbf{Z}}^{\log}$) is equal to the chosen $(\mathbf{Q}_1)_{\mathbf{Z}/p^2\mathbf{Z}}$ modulo $p \cdot F^1$. This gives us a natural choice of

$$(\mathbf{Q}_1)_{\mathbf{Z}/p^2\mathbf{Z}} \stackrel{\text{def}}{=} \mathbf{F}^*((\mathbf{P}_0)_{\mathbf{Z}/p^3\mathbf{Z}})_{\mathbf{Z}/p^2\mathbf{Z}} \quad \text{modulo } p^2$$

(i.e., not just modulo $p \cdot F^1$). Then we may choose a lifting of $(\mathbf{Q}_1)_{\mathbf{Z}/p^3\mathbf{Z}}$ modulo $p^2 \cdot F^1$. Thus, since the morphism $(\Theta_1)_S$ (cf. the discussion of §3.2) is an *isomorphism*, there is a unique lifting $\theta_{\mathbf{Z}/p^3\mathbf{Z}}^{\log} : S_{\mathbf{Z}/p^3\mathbf{Z}}^{\log} \rightarrow T_{\mathbf{Z}/p^3\mathbf{Z}}^{\log}$ whose corresponding $\mathbf{F}^*((\mathbf{Q}_1)_{\mathbf{Z}/p^3\mathbf{Z}})$ is spk-indigenous modulo p^2 (cf. the above discussion). This gives us a natural choice for $(\mathbf{P}_0)_{\mathbf{Z}/p^2\mathbf{Z}} \stackrel{\text{def}}{=} \mathbf{F}^*((\mathbf{Q}_1)_{\mathbf{Z}/p^3\mathbf{Z}})_{\mathbf{Z}/p^2\mathbf{Z}}$. Continuing in this fashion thus gives us *unique* liftings

$$\theta_{\mathbf{Z}_p}^{\log} : S_{\mathbf{Z}_p}^{\log} \rightarrow T_{\mathbf{Z}_p}^{\log}; \quad \phi_{\mathbf{Z}_p}^{\log} : T_{\mathbf{Z}_p}^{\log} \rightarrow S_{\mathbf{Z}_p}^{\log}$$

for which the corresponding renormalized Frobenius pull-backs of certain bundles possess certain properties. Note that this *uniqueness*, together with the rigidification of $T_{\mathbf{Z}_p}^{\log}$ provided by the deformation $(\mathbf{Q}_1)_{\mathbf{Z}_p}$ (a crystal on $\text{Crys}(X_T^{\log}/T_{\mathbf{Z}_p}^{\log})$), allow us to glue together such $\theta_{\mathbf{Z}_p}^{\log}$, $\phi_{\mathbf{Z}_p}^{\log}$ after they have been constructed étale locally on $T_{\mathbf{Z}_p}$. Moreover, by the uniqueness part of Theorem 1.8, we have that

$$\Phi^{\log} = \phi_{\mathbf{Z}_p}^{\log} \circ \theta_{\mathbf{Z}_p}^{\log} : S_{\mathbf{Z}_p}^{\log} \rightarrow S_{\mathbf{Z}_p}^{\log}$$

That is to say, we have proven the following result:

Theorem 3.8. *Let Π be a spiked VF-pattern of period $\varpi = 2$. Let $S_{\mathbf{Z}_p} \rightarrow (\overline{\mathcal{N}}_{g,r}^{\Pi,s})_{\mathbf{Z}_p}^{\text{ord}}$ be an étale morphism of p-adic formal stacks that maps into the very ordinary locus (cf. Definition 3.6). (In fact, it suffices to make the weaker assumption that it maps into the locus where condition (a.) of Definition 3.6 is satisfied.) Equip $S_{\mathbf{Z}_p}$ with the log structure pulled back from $\overline{\mathcal{M}}_{g,r}^{\log}$. Let T^{\log} be as in Lemma 3.3. Then there exists a natural \mathbf{Z}_p -flat p-adic formal log stack $T_{\mathbf{Z}_p}^{\log}$, together with canonical Frobenius liftings*

$$\theta_{\mathbf{Z}_p}^{\log} : S_{\mathbf{Z}_p}^{\log} \rightarrow T_{\mathbf{Z}_p}^{\log}$$

and

$$\phi_{\mathbf{Z}_p}^{\log} : T_{\mathbf{Z}_p}^{\log} \rightarrow S_{\mathbf{Z}_p}^{\log}$$

of $\theta^{\log} \stackrel{\text{def}}{=} \Phi^{\text{Hs},\log} : S^{\log} \rightarrow T^{\log}$ and $\phi^{\log} \stackrel{\text{def}}{=} \Phi_S^{\log} \circ \Phi^{\text{spk},\log} : T^{\log} \rightarrow S^{\log}$ (where $\Phi^{\text{Hs},\log}$ and $\Phi^{\text{spk},\log}$ are as in Lemma 3.3). Moreover, the composite of $\theta_{\mathbf{Z}_p}^{\log}$ and $\phi_{\mathbf{Z}_p}^{\log}$ is equal to the canonical Frobenius lifting

$$\Phi^{\log} : S_{\mathbf{Z}_p}^{\log} \rightarrow S_{\mathbf{Z}_p}^{\log}$$

of Theorem 1.8. Finally, the derivatives of these Frobenius liftings have the following property: Let us denote by $\Omega'_{T_{\mathbf{Z}_p}^{\log}} \subseteq \Omega_{T_{\mathbf{Z}_p}^{\log}}$ the subsheaf of sections whose reductions modulo p are annihilated by $\Theta_T^{\text{spk}} = \text{Ker}(\delta^{\text{spk}}) (\subseteq \Theta_{T^{\log}} = (\Omega_{T^{\log}})^{\vee})$. Then (after dividing by p in both cases), $d\theta_{\mathbf{Z}_p}^{\log}$ induces an isomorphism

$$\theta_{\mathbf{Z}_p}^* \Omega'_{T_{\mathbf{Z}_p}^{\log}} \rightarrow \Omega_{S_{\mathbf{Z}_p}^{\log}}$$

while $d\phi_{\mathbf{Z}_p}^{\log}$ induces an isomorphism

$$\phi_{\mathbf{Z}_p}^* \Omega_{S^{\log}_{\mathbf{Z}_p}} \rightarrow \Omega_{T^{\log}_{\mathbf{Z}_p}}$$


We shall refer to this property of the derivatives by saying that the system of Frobenius liftings $\{\theta_{\mathbf{Z}_p}^{\log}, \phi_{\mathbf{Z}_p}^{\log}\}$ is Π -ordinary.

Proof. It remains only to prove the assertion concerning the derivatives. But this follows by observing that if one reduces modulo p the two morphisms in question (and then pulls back to $\mathcal{S} \stackrel{\text{def}}{=} (S^{\log})^{\text{pro}}$), one obtains the inverses to the duals of the isomorphisms $(\Phi_1)_{\mathcal{S}}$ and $(\Phi_0)_{\mathcal{S}}$ discussed at the beginning of §3.2. Indeed, this relationship between the derivatives discussed here and the isomorphisms $(\Phi_1)_{\mathcal{S}}$ and $(\Phi_0)_{\mathcal{S}}$ of §3.2 is a tautology of the same type as that discussed in Lemma 2.3. \circ


Remark. It would be nice if we could prove the existence of such a decomposition for every spiked canonical system of Frobenius liftings (as in Theorem 1.8). Unfortunately, however, without the rather strong assumptions made here (e.g., that the period is 2; condition (a.) of Definition 3.6), this does not seem to be possible. For instance, when the period is > 2 , it is not even clear how to construct nice intermediate “ T ”’s at each step. This is why we had to settle for composite Frobenius liftings in Theorem 1.8 that “skip” the spiked steps, and it is also why we had to restrict ourselves to the binary case in §2.

Pictorial Appendix

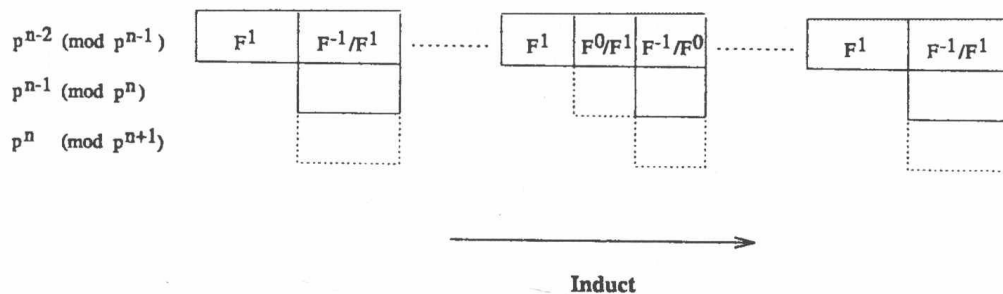
In this Appendix, we give an illustration of what is going on in the inductive lifting procedure of Theorem 1.8. First we give an illustration of the general situation. Let (i, j) be a Π -ind-adjacent pair (so $i < j$). Then the various portions of liftings of the i^{th} (respectively, j^{th}) link bundle – which will be indigenous – are shown on the right (respectively, left). The various portions of liftings of the typical n^{th} link bundle (where $i < n < j$) are shown in the middle.



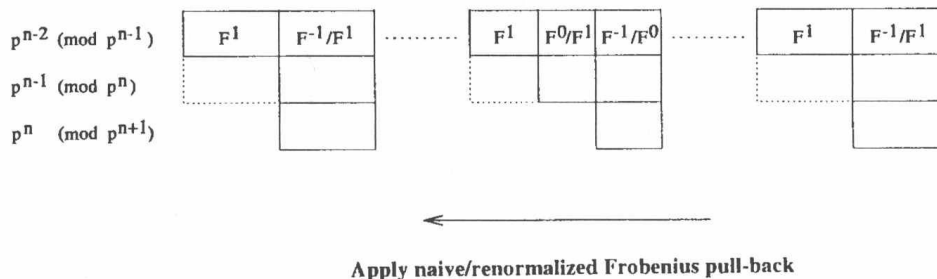
Next, we give an illustration of what has been determined at the beginning of the n^{th} induction step:



Next, we give an illustration of what is going on in the first substep of the n^{th} induction step; the portions just determined are shown in dotted lines:



Finally, we give an illustration of what is going on in the last substep of the n^{th} induction step; the portions just determined are shown in dotted lines:



Chapter VIII: The Geometrization of Binary-Ordinary Frobenius Liftings

§0. Introduction

In this Chapter, we attempt to generalize §1 of Chapter III of [Mzk1] (where we showed how “(classical) ordinary Frobenius liftings” give rise to certain canonical coordinates – cf. Theorem 0.3 of the Introduction to the present work for a synopsis of this portion of [Mzk1]) to the present context, i.e., systems of Frobenius liftings as in Chapter VII, Theorem 1.8. Unfortunately, because of the complexity of an arbitrary system as in Chapter VII, Theorem 1.8, we do not succeed in general, and our best results are obtained only in the *binary* (in this Chapter) and *very ordinary spiked* (in the next Chapter) cases.

Unlike the *classical ordinary case* treated in Chapter III of [Mzk1], here we obtain much more complicated and interesting geometries. In the classical ordinary case, the geometry obtained was based on a local isomorphism to (products of copies of) \widehat{G}_m . Here, however, we obtain (for instance, in the case when the VF-pattern Π is of pure tone ϖ) *Lubin-Tate geometries* based on local isomorphisms to various Lubin-Tate formal groups. Moreover, in the *anabelian* case (when the VF-pattern Π is such that there are several active integers in a period), one obtains certain *noncommutative geometries*, which PD-locally look like products of “abelian geometries” (as in the classical ordinary and Lubin-Tate cases), but in which the geodesic directions are tangled up in some sort of noncommuting fashion. In §1, 2, we discuss the *Galois representations, geometries, and uniformizations* (Theorems 2.12, 2.15, and 2.17) associated to certain types of (namely, “*binary-ordinary*”) systems of Frobenius liftings in a rather general context (which has nothing to do with curves or their moduli), and in §3, we explain the relationship between these general geometric ideas and the theory (involving curves (Theorem 3.2) and their moduli (Theorem 3.1)) developed in earlier chapters.

In particular, we obtain, for each binary VF-pattern Π , a new (i.e., generalizing that of [Mzk1]) canonical uniformization (étale locally) of $(\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{Z}_p}$ and the universal curve $\mathcal{C}_{\mathbf{Z}_p}^{\log} \rightarrow (\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{Z}_p}$.

Indeed, one way to view the last three chapters (i.e., Chapters VIII, IX, and X) of this work in the context of the entire work is the following:

Up till now, we have been concerned with constructing various canonical, Frobenius invariant, p -adic differential equations on hyperbolic curves and their moduli. The last three chapters then are devoted to the task of integrating these differential equations to obtain functions which serve to uniformize hyperbolic curves and their moduli.

Finally, for illustrations of the author's conception of the new sorts of geometries obtained, we refer to the Pictorial Appendix (as well as §1.6 of the Introduction to the present work).

§1. The General Framework

In this §, we make some basic definitions and discuss what this Chapter will be about.

§1.1. Canonical Points

Let k be a perfect field of odd characteristic p . Let $A = W(k)$, the ring of Witt vectors with coefficients in k ; let K be its quotient field. Let S be a formally smooth, geometrically connected p -adic formal scheme over A of constant relative dimension d . Let S^{\log} be a log formal scheme whose underlying formal scheme is S and whose log structure is given by a relative divisor with normal crossings $D \subseteq S$ over A . Let $\Phi_A : A \rightarrow A$ be the Frobenius morphism on A . Let us denote the result of base changing by Φ_A by means of a superscripted " F ." Let $n \geq 1$ be an integer. Let us assume that for $i = 1, \dots, n$, we are given a morphism

$$\Phi_i^{\log} : S^{\log} \rightarrow S^{\log}$$

whose reduction modulo p is the λ_i^{th} (where $\lambda_i \geq 1$ is an integer) power of the usual Frobenius morphism in characteristic p . Let Λ be the ordered set $\{\lambda_1, \dots, \lambda_n\}$. Let

$$\varpi = \sum_{i=1}^n \lambda_i$$

Definition 1.1. We shall refer to the data $\{\Phi_i^{\log}\}_{i=1,\dots,n}$ as a *system of Frobenius liftings*, or *Frobenius system (of multi-order Λ and period ϖ)*.

Let σ be a permutation of $\{1, \dots, n\}$ such that there exists a $\mu \in \mathbf{Z}$ such that $\sigma(i) \equiv i + \mu \pmod{n}$. We shall call such a σ a *shifting permutation*. Let

$$\Phi_{\sigma}^{\log} : S^{\log} \rightarrow S^{\log}$$

be the composite $\Phi_{\sigma(1)}^{\log} \circ \dots \circ \Phi_{\sigma(n)}^{\log}$.

Definition 1.2. We shall refer to Φ_{σ}^{\log} as the σ -*shifted total Frobenius lifting* associated to the above Frobenius system.

As in the discussion preceding Theorem 1.9 of Chapter VII, we have a natural commutative diagram:

$$\begin{array}{ccc} W((S^{\log})_{\mathbf{F}_p}^{\text{pro}}) & \xrightarrow{\phi_{\sigma}} & S^{\log} \\ \downarrow & & \downarrow \Phi_{\sigma}^{\log} \\ W((S^{\log})_{\mathbf{F}_p}^{\text{pro}}) & \xrightarrow{\phi_{\sigma}} & S^{\log} \end{array}$$

where the vertical morphism on the left is the ϖ^{th} power of the canonical Frobenius on the Witt vectors, and the horizontal morphisms are both equal to a certain morphism $\phi_{\sigma} : W((S^{\log})_{\mathbf{F}_p}^{\text{pro}}) \rightarrow S^{\log}$ canonically defined by Φ_{σ}^{\log} (and subject to the condition that the reduction modulo p of the morphism ϕ_{σ} is the natural morphism $(S^{\log})_{\mathbf{F}_p}^{\text{pro}} \rightarrow S_{\mathbf{F}_p}^{\log}$). In particular, note that for any perfect log scheme T^{\log} (of characteristic p) and any morphism $\tau : T^{\log} \rightarrow S_{\mathbf{F}_p}^{\log}$, there exists a unique morphism $\tau_{\sigma} : W(T^{\log}) \rightarrow S^{\log}$ (i.e., a $W(T^{\log})$ -valued point of S^{\log}) such that the following diagram commutes

$$\begin{array}{ccc} W(T^{\log}) & \xrightarrow{\tau_{\sigma}} & S^{\log} \\ \downarrow & & \downarrow \Phi_{\sigma}^{\log} \\ W(T^{\log}) & \xrightarrow{\tau_{\sigma}} & S^{\log} \end{array}$$

(where the vertical morphism on the left is the ϖ^{th} power of the canonical Frobenius on the Witt vectors). Then we make the following definition (cf. [Mzk1], Chapter III, Definition 1.9):

Definition 1.3. We shall refer to $\tau_\sigma \in S^{\log}(W(T^{\log}))$ as the σ -canonical lifting of $\tau \in S^{\log}(T^{\log})$. We shall refer to a point of $S^{\log}(W(T^{\log}))$ which is the σ -canonical lifting of some point of $S^{\log}(T^{\log})$ as a σ -canonical ($W(T^{\log})$ -valued) point of S^{\log} (relative to the given Frobenius system).

Remark. Note that a point which is σ -canonical for the system will not, in general, be σ' -canonical for $\sigma' \neq \sigma$.

Now fix a point $s \in S^{\log}(K)$ at which the log structure is trivial. Let

$$\Pi_{S^{\log}} \stackrel{\text{def}}{=} \pi_1(S_K^{\log}, s)$$

be the (logarithmic) fundamental group of S_K^{\log} . Note that for each $m \geq 1$, we have a log étale covering of S_K^{\log} given by taking the m^{th} iterate of Φ_σ^{\log} and tensoring with K :

$$(\Phi_\sigma^{\log})_K^m : S_K^{\log} \rightarrow S_K^{\log}$$

Let us call the corresponding $\Pi_{S^{\log}}$ -set (i.e., profinite set with continuous $\Pi_{S^{\log}}$ action) $\mathcal{U}_{\sigma,m}$. Observe that the underlying set of $\mathcal{U}_{\sigma,m}$ has $p^{\varpi \cdot d \cdot m}$ elements. As m varies, we obtain $\Pi_{S^{\log}}$ -equivariant surjections

$$\mathcal{U}_{\sigma,m+1} \rightarrow \mathcal{U}_{\sigma,m}$$

Thus, the $\mathcal{U}_{\sigma,m}$'s form an inverse system of $\Pi_{S^{\log}}$ -sets. Let

$$\mathcal{U}_\sigma$$

be the inverse limit of the $\mathcal{U}_{\sigma,m}$'s. Thus, \mathcal{U}_σ will be a compact (profinite) $\Pi_{S^{\log}}$ -set. Let

$$\mathcal{U} \stackrel{\text{def}}{=} \prod_{\sigma} \mathcal{U}_\sigma$$

be the product of the \mathcal{U}_σ over all shifting permutations σ . Similarly, we write \mathcal{U}_m for the product of the $\mathcal{U}_{\sigma,m}$ over all σ .

Definition 1.4. We shall refer to the $\Pi_{S^{\log}}$ -set \mathcal{U} as the set-theoretic canonical Galois representation associated to the above Frobenius system.

§1.2. The Meaning of “Geometrization”

We are now in a position to discuss what we mean by the word “geometrization” in the title of this Chapter. Let

$$\mathcal{D}_S$$

be the PD-envelope of the diagonal in $S^{\log} \times_A S^{\log}$. By regarding \mathcal{D}_S as an $\mathcal{O}_{S^{\log}}$ -algebra from the right, one sees easily that \mathcal{D}_S admits a *connection* (on S^{\log} over A), as well as a *Hodge filtration* defined by divided powers of the ideal defining the diagonal embedding (cf., e.g., [Mzk1], Chapter V, §1). Moreover, Φ_{id}^{\log} induces a *Frobenius action* on \mathcal{D}_S . Thus, *morally*, \mathcal{D}_S with these added structures should define a sort of *infinite \mathcal{MF}^{∇} -object* (i.e., an infinite rank version of the (finite rank) objects in the category \mathcal{MF}^{∇} of [Falt1], §2), hence should give rise to a *continuous Galois representation* on some algebra that looks roughly like the p -adic completion of $\mathbf{Z}_p[t_1, \dots, t_{dn}]$ (where the t_i ’s are indeterminates). Moreover, the induced Galois representation on the \mathbf{Z}_p -valued points of this algebra should be naturally isomorphic to \mathcal{U} .

Unfortunately, there are several technical problems here. Of these, the most fundamental is that in general, one does not have a good theory of such “infinite \mathcal{MF}^{∇} -objects.” In fact, the theory of [Falt1], §2, only works well when the length of the Hodge filtration is $\leq p - 1$. In [Mzk1], Chapter V, §1, we got around this problem by observing that there exists a *sub- \mathcal{MF}^{∇} -object inside \mathcal{D}_S whose Hodge filtration has finite length* (actually, the length was 2) and which (essentially) generates \mathcal{D}_S as an \mathcal{O}_S -algebra. This allowed us to actually associate a Galois representation to the \mathcal{MF}^{∇} -object \mathcal{D}_S in a natural way.

Another way to think of this finite sub- \mathcal{MF}^{∇} -object inside \mathcal{D}_S is that it corresponds to the *geometry* associated to a (classical) ordinary Frobenius lifting. That is to say, as was seen in [Mzk1], Chapter III, §1, a (classical) ordinary Frobenius lifting on S^{\log} induces a sort of *affine geometry* on S^{\log} , and this finite sub- \mathcal{MF}^{∇} -object corresponds (relative to this geometry) to the functions that “respect this geometry.” Indeed, from the “Clifford-Klein point of view,” a *geometry* (on a real differentiable manifold M) is precisely an action of a Lie group on the universal cover \widetilde{M} of the manifold (satisfying certain conditions). By looking at the action of the Lie group on the C^{∞} -functions $C^{\infty}(\widetilde{M})$ on this universal cover, one sees (typically) that there exist *finite-dimensional sub-representations* of $C^{\infty}(\widetilde{M})$. Moreover, such a finite-dimensional sub-representation which is big enough to define an embedding of \widetilde{M} in fact determines the *geometry on \widetilde{M}* (and hence on M). Thus,

Suppose that we define a p -adic geometry (on S^{\log} and relative to the given Frobenius system) as a sub- \mathcal{MF}^{∇} -object of \mathcal{D}_S which

“essentially” generates \mathcal{D}_S as an \mathcal{O}_S -algebra (and which, ideally, is of finite rank, and minimal among such generating sub-objects). (Here, we regard \mathcal{D}_S as an “infinite \mathcal{MF}^∇ -object” by means of the Frobenius action defined by Φ_{id}^{\log} .) Then one sees that there is a direct analogy between real differentiable, Clifford-Klein geometries and p -adic geometries (as just defined).

Typically, in the real differentiable case, the geometry can be given locally in terms of a *metric* (on the tangent bundle) which defines (at least locally) *normal coordinates*. These normal coordinates can be used (at least locally) as the “geometrizing” finite-dimensional sub-representation in the above discussion. We shall refer to geometries that can be defined by such metrics as (locally) *affine*. One may also think of “locally affine geometries” as being geometries for which:

- (1) (at least locally) the dimension of the “geometrizing” finite-dimensional sub-representation is the *same* as the dimension of M ; or, alternatively,
- (2) the differential equation for a geodesic $\gamma(t)$ (which states essentially that “the velocity is constant”) depends only on the *first derivative* of $\gamma'(t)$ (this is also reflected in the fact that the metric is on the *tangent bundle*).

(Note that “locally affine” in the real differentiable case corresponds to “affine” in the p -adic case, since in the p -adic case, one is ultimately only concerned with \mathcal{D}_S , which consists of *local* functions.) In general, however, one may want to consider more complicated geometries, which are *not locally affine*, i.e., geometries in which the differential equation for a geodesic involves higher derivatives of $\gamma'(t)$.

As stated above, we do not succeed in this Chapter in defining p -adic geometries for arbitrary Frobenius systems (or even arbitrary Frobenius systems of the type that appear in Theorem 1.8 of Chapter VII). However, at least in the *binary case*, we succeed in associating a *natural* (in general, “nonabelian” – see §2.6 below for more on the meaning of this) *affine geometry to the canonical Frobenius system*. On the other hand, even in the *simplest of the spiked cases* (i.e., the very ordinary spiked case), we are unable to show the existence of a natural *affine geometry* associated to the given Frobenius lifting. That is to say, in the very ordinary spiked case, the natural geometry that we will associate to the given Frobenius lifting (in Chapter IX) *will not be affine*. Thus, one must expect that for an arbitrary spiked canonical Frobenius system, it will be all the more difficult (if not impossible) to associate to the system an affine geometry.

§2. The Binary Case

§2.1. The Associated Differential Formal Group

Let us maintain the notation of §1.1. Let $\mathcal{O}_\varpi = W(\mathbf{F}_{p^\varpi})$; $K_\varpi = \mathcal{O}_\varpi \otimes_{\mathbf{Z}_p} \mathbf{Q}_p$. Let us assume henceforth that $\mathbf{F}_{p^\varpi} \subseteq k$. Note that, by dividing by p , each Φ_i^{\log} induces a morphism

$$\Omega_{\Phi_i} \stackrel{\text{def}}{=} \frac{1}{p} \cdot d\Phi_i^{\log} : \Phi_i^* \Omega_{S^{\log}/A} \rightarrow \Omega_{S^{\log}/A}$$

Now we make the following

Definition 2.1. We shall say that the Frobenius system in question is *binary-ordinary* if the following conditions are satisfied:

- (1) all the Ω_{Φ_i} 's are isomorphisms;
- (2) either all the λ_i 's are 1, or the log structure on S^{\log} is trivial.

When we want to make clear that we mean “ordinary” in the sense of [Mzk1], we shall use the term *classical ordinary*.

Let us assume that our Frobenius system is *binary-ordinary*. Then (just as in the classical ordinary case) we would like to associate various p -divisible groups to our Frobenius system. Let σ be a shifting permutation. Then we define

$$\Omega_{\Phi_\sigma} \stackrel{\text{def}}{=} \frac{1}{p^n} \cdot d\Phi_{\sigma(n)}^{\log} \circ \dots \circ d\Phi_{\sigma(1)}^{\log} : \Phi_\sigma^* \Omega_{S^{\log}/A} \rightarrow \Omega_{S^{\log}/A}$$

By the assumption of binary-ordinariness, Ω_{Φ_σ} is an *isomorphism*. Thus, by taking Ω_{Φ_σ} -invariants, we obtain an *étale local system*

$$\Omega_{\Phi_\sigma}^{\text{ét}}$$

in free \mathcal{O}_ϖ -modules of rank $d = \dim_A(S)$ on S . Moreover, by taking the elements of $\Omega_{\Phi_\sigma}^{\text{ét}}$ in Ω_{Φ_σ} to be horizontal, we obtain a *connection*

$$\nabla_{\Omega_\sigma}$$

on $\Omega_{S^{\log}/A}$.

Definition 2.2. We shall refer to $\Omega_{\Phi_\sigma}^{\text{et}}$ (respectively, its dual $\Theta_{\Phi_\sigma}^{\text{et}}$) as the σ -canonical differential (respectively, tangential) local system on S associated to the given Frobenius system. We shall refer to ∇_{Ω_σ} as the σ -canonical connection on $\Omega_{S^{\log}/A}$.

Note that if $\sigma(1) = i$ and $\sigma'(1) = i + 1$, then applying Ω_{Φ_i} to $\Omega_{\Phi_\sigma}^{\text{et}}$ maps $\Omega_{\Phi_\sigma}^{\text{et}}$ bijectively onto $\Omega_{\Phi_{\sigma'}}^{\text{et}}$. If we take into account the \mathcal{O}_ϖ -module structures, then this bijection gives a natural \mathcal{O}_ϖ -linear isomorphism

$$\Omega_{\Phi_\sigma}^{\text{et}} \otimes_{\mathcal{O}_\varpi, \Phi_{\mathcal{O}_\varpi}^{\lambda_i}} \mathcal{O}_\varpi \cong \Omega_{\Phi_{\sigma'}}^{\text{et}},$$

where $\Phi_{\mathcal{O}_\varpi}$ is the Frobenius on \mathcal{O}_ϖ .

Unlike the classical ordinary case, however, in fact, there exists a much finer structure on $\Omega_{S^{\log}/A}$, defined as follows. For $i = 1, \dots, n$, let $M_i \stackrel{\text{def}}{=} \Omega_{S^{\log}/A}$. Endow M_i with the connection ∇_{M_i} its σ -canonical connection (Definition 2.2) for the σ for which $\sigma(1) = i$. Thus,

$$\mathcal{M}_i \stackrel{\text{def}}{=} (M_i, \nabla_{M_i})$$

forms a crystal on $\text{Crys}(S \otimes \mathbf{F}_p/A)$, so we can pull it back by the Frobenius $\Phi_{S_{\mathbf{F}_p}}$ on $S \otimes \mathbf{F}_p$. Now let $\mathcal{M}_{i,j} \stackrel{\text{def}}{=} (\Phi_{S_{\mathbf{F}_p}}^*)^j \mathcal{M}_i$, and let

$$\mathcal{M} \stackrel{\text{def}}{=} \bigoplus_{i=1}^n \left(\bigoplus_{j=0}^{\lambda_i-1} \mathcal{M}_{i,j} \right)$$

Thus, if we write $\mathcal{M} = (M, \nabla_M)$, then M is a vector bundle of rank $\varpi \cdot d$ on S .

Next, we would like to define a Hodge filtration on M : We let

$$F^0(M) \stackrel{\text{def}}{=} M; \quad F^1(M) \stackrel{\text{def}}{=} \bigoplus_{i=1}^n M_{i,0}; \quad F^2(M) \stackrel{\text{def}}{=} 0$$

Finally, we would like to define a Frobenius action, as follows: Suppose that $j = \lambda_i - 1$; then on $\mathcal{M}_{i,j}$, if $j > 0$ (respectively, $j = 0$), the Frobenius action is given by the isomorphism (respectively, p times the isomorphism) $\Phi_{S_{\mathbf{F}_p}}^* \mathcal{M}_{i,j} \cong \mathcal{M}_{i+1,0}$ induced by Ω_{Φ_i} . (Here we regard $\mathcal{M}_{n+1,0}$ as $\mathcal{M}_{1,0}$.) Now suppose that $j < \lambda_i - 1$; then on $\mathcal{M}_{i,j}$, if $j > 0$ (respectively, $j = 0$), the Frobenius action is given by the identity (respectively, p times the identity) $\Phi_{S_{\mathbf{F}_p}}^* \mathcal{M}_{i,j} = \mathcal{M}_{i,j+1}$. One sees easily that with this Frobenius action Φ_M , we obtain an \mathcal{MF}^∇ -object

$$(M, \nabla_M, F(M), \Phi_M)$$

Moreover, it follows from the way we defined the Hodge filtration that this \mathcal{MF}^∇ -object defines a p -divisible group G_{Ω_Φ} on S^{\log} (cf. [Falt1], §7, Theorem 7.1). (Note that here, we use the *contravariant* correspondence between p -divisible groups and their Dieudonné crystals.)

Definition 2.3. We shall refer to G_{Ω_Φ} as the *differential p -divisible group on S^{\log} associated to the given Frobenius system*.

Next, let us note that G_{Ω_Φ} has a *natural \mathcal{O}_ϖ -action*, given as follows. Given an element $\alpha \in \mathcal{O}_\varpi \subseteq A$, we let α act on $M_{i,j}$ as $\Phi_A^N(\alpha)$ (via the usual A -module structure on $M_{i,j}$), where $N = j + \sum_{a=1}^{i-1} \lambda_a$. This gives an action of \mathcal{O}_ϖ on M that commutes with Φ_M and ∇_M , and preserves the Hodge filtration. Thus, *we obtain an action of \mathcal{O}_ϖ on G_{Ω_Φ}* .

On the other hand, let us define

$$G_\Lambda$$

over $\mathrm{Spec}(A)$ (in a fashion analogous to the definition of G_{Ω_Φ}) as follows: For $i = 1, \dots, n$; $j = 0, \dots, \lambda_i - 1$, let

$$(M_\Lambda)_{i,j} \stackrel{\mathrm{def}}{=} A$$

Let

$$M_\Lambda \stackrel{\mathrm{def}}{=} \bigoplus_{i=1}^n \left(\bigoplus_{j=0}^{\lambda_i-1} (M_\Lambda)_{i,j} \right)$$

Let

$$F^0(M_\Lambda) \stackrel{\mathrm{def}}{=} M_\Lambda; \quad F^1(M_\Lambda) \stackrel{\mathrm{def}}{=} \bigoplus_{i=1}^n (M_\Lambda)_{i,0}; \quad F^2(M_\Lambda) \stackrel{\mathrm{def}}{=} 0$$

Finally, we define a *Frobenius action*, as follows: Suppose that $j = \lambda_i - 1$; then on $(M_\Lambda)_{i,j}$, if $j > 0$ (respectively, $j = 0$), the Frobenius action is given by the identity (respectively, p times the identity) $\Phi_A^*(M_\Lambda)_{i,j} = A = (M_\Lambda)_{i+1,0}$. Now suppose that $j < \lambda_i - 1$; then on $(M_\Lambda)_{i,j}$, if $j > 0$ (respectively, $j = 0$), the Frobenius action is given by the identity (respectively, p times the identity) $\Phi_A^*(M_\Lambda)_{i,j} = A = (M_\Lambda)_{i,j+1}$. One sees easily that with this Frobenius action, we obtain an \mathcal{MF}^∇ -object on

$\mathrm{Spec}(A)$, and hence a p -divisible group G_Λ (on $\mathrm{Spec}(A)$). Naturally, G_Λ also admits a natural \mathcal{O}_ϖ -action. Now we are ready for the following

Proposition 2.4. *We have a natural isomorphism*

$$G_{\Omega_\Phi} \cong \Theta_{\Phi, \mathrm{id}}^{\mathrm{et}} \otimes_{\mathcal{O}_\varpi}^{\mathrm{gp}} (G_\Lambda)|_S$$

Here, (as in the rest of this chapter) we use the symbol “ \otimes^{gp} ” to denote the tensor product of “ \mathcal{O}_ϖ -module schemes” (i.e., schemes that take values in the category of \mathcal{O}_ϖ -modules).

Proof. Indeed, this is immediate from the definitions. \circ

Since the p -divisible group G_Λ on $\mathrm{Spec}(A)$ will play a central role in this chapter, let us examine it in greater detail. The following observations are immediate:

- (1) G_Λ is naturally defined as a p -divisible group (respectively, p -divisible group with \mathcal{O}_ϖ -action) over \mathbf{Z}_p (respectively, \mathcal{O}_ϖ). Indeed, this follows immediately from looking at the defining \mathcal{MF}^∇ -object.
- (2) Since G_Λ has no étale part, it defines a *connected formal group*

$$\mathcal{G}_\Lambda \rightarrow \mathrm{Spec}(A)$$

If we let $\mathcal{G}_\Lambda[p^i]$ be the kernel of multiplication by p^i on \mathcal{G}_Λ , then one sees easily that one can recover G_Λ from \mathcal{G}_Λ as the inductive limit of the $\mathcal{G}_\Lambda[p^i]$.

- (3) By the elementary theory of Dieudonné modules of p -divisible groups (cf., e.g., [Font2]), the fact that $\mathrm{rank}_A(F^1(M_\Lambda)) = n$ implies that \mathcal{G}_Λ is formally smooth over A of relative dimension n .
- (4) The sheaf of invariant differentials on \mathcal{G}_Λ can be naturally identified with

$$F^1(M_\Lambda) = A^n$$

(i.e., a direct sum of n copies of A). Thus, there are canonical differentials $\omega_{\Lambda, \sigma}$ (where σ is a shifting permutation such that $\sigma(1) = i$) on \mathcal{G}_Λ given by the elements of

$F^1(M_\Lambda)$ corresponding to $(0, \dots, 0, 1, 0, \dots, 0)$ (where the 1 is in the i^{th} place).

- (5) Let $\Theta_K(\mathcal{G}_\Lambda)$ be the *tangent space* (over K) to the identity of \mathcal{G}_Λ . By (4) above, we have a natural identification of $\Theta(\mathcal{G}_\Lambda)$ with K^n . Moreover, as for any formal group, we have a *logarithm map*

$$\log_\Lambda : \mathcal{G}_\Lambda(A) \rightarrow \Theta_K(\mathcal{G}_\Lambda) = K^n$$

defined by the fact that in characteristic zero, all formal groups are isomorphic to the additive group.

Now let R be a *local artinian A -algebra such that the structure morphism $A \rightarrow R$ induces an isomorphism of residue fields*. Then we have the following important technical result:

Lemma 2.5. *Suppose that we have an exact sequence of formal groups on the finite flat site of $\text{Spec}(R)$*

$$0 \rightarrow \mathcal{G}_\Lambda|_{\text{Spec}(R)} \rightarrow \mathcal{H} \rightarrow \mathbf{Z} \rightarrow 0$$

(where “ \mathbf{Z} ” denotes the constant formal group with fiber \mathbf{Z} over $\text{Spec}(R)$) that splits over $\text{Spec}(k) \subseteq \text{Spec}(R)$. Then the sequence splits over $\text{Spec}(R)$.

Proof. This follows immediately from the formal smoothness of \mathcal{G}_Λ over A : Indeed, a splitting over $\text{Spec}(k)$ consists of a point $\alpha_k \in \mathcal{H}(k)$ that maps to $1 \in \mathbf{Z}$. Since \mathcal{G}_Λ is formally smooth over A , it follows that \mathcal{H} is smooth over R , so α_k lifts to a point $\alpha_R \in \mathcal{H}(R)$. Clearly, α_R defines a splitting of the sequence. \square

The purpose of the above Lemma is that it allows us to do “Kummer theory” with respect to G_Λ . Indeed, suppose that R is as above, and suppose that we have an exact sequence

$$0 \rightarrow G_\Lambda|_{\text{Spec}(R)} \rightarrow H \rightarrow \mathbf{Q}_p/\mathbf{Z}_p \rightarrow 0$$

on the finite flat site of $\text{Spec}(R)$. Since k is *perfect*, it follows immediately (for instance, by considering the direct sum decomposition of the Dieudonné module of H into its slope zero and slope > 0 components) that *this sequence splits* over $\text{Spec}(k)$. On the other hand, the kernel of p^N on G_Λ is the same as the kernel of p^N on \mathcal{G}_Λ . Thus, by looking at the long exact cohomology sequence associated to the exact sequence

$$0 \longrightarrow \mathcal{G}_\Lambda[p^N]|_{\mathrm{Spec}(R)} \longrightarrow \mathcal{G}_\Lambda|_{\mathrm{Spec}(R)} \xrightarrow{p^N} \mathcal{G}_\Lambda|_{\mathrm{Spec}(R)} \longrightarrow 0$$

on the finite flat site of $\mathrm{Spec}(R)$ and applying Lemma 2.5 above, we obtain (as in Kummer theory) that the above extension involving H corresponds to a class

$$\eta_H \in \varprojlim \mathcal{G}_\Lambda(R)/p^N \cdot \mathcal{G}_\Lambda(R) = \mathcal{G}_\Lambda(R)$$

where the inverse limit is over integers $N \geq 1$, and the last equality follows from the fact that for N large, p^N annihilates $\mathcal{G}_\Lambda(R)$. Thus, we have the following:

Proposition 2.6. (Kummer Theory for \mathcal{G}_Λ) *Let R be a local artinian A -algebra such that the structure morphism $A \rightarrow R$ induces an isomorphism of residue fields. Then there is a natural bijective correspondence between extensions*

$$0 \rightarrow G_\Lambda|_{\mathrm{Spec}(R)} \rightarrow H \rightarrow \mathbf{Q}_p/\mathbf{Z}_p \rightarrow 0$$

on the finite flat site of $\mathrm{Spec}(R)$, and elements

$$\eta_H \in \mathcal{G}_\Lambda(R)$$

Moreover, this correspondence is defined by considering the extension (as above) obtained by dividing η_H by powers of p .

Remark. Note that Proposition 2.6 is a “non-logarithmic version” of Kummer theory. When all the λ_i ’s are 1, there is also a logarithmic version, whose formulation we leave to the reader.

We end this subsection with some *examples of \mathcal{G}_Λ* :

Example 2.7. Suppose that $n = 1$. Then \mathcal{G}_Λ is determined by $\lambda_1 = \varpi$. If $\varpi = 1$, then \mathcal{G}_Λ is the completion $\widehat{\mathbf{G}}_{\mathfrak{m}}$ of the multiplicative group $\mathbf{G}_{\mathfrak{m}}$ at the identity. More generally, one sees easily that (for arbitrary ϖ) \mathcal{G}_Λ is the *Lubin-Tate formal group associated to \mathcal{O}_ϖ* . Indeed, this follows from the fact that the Lubin-Tate group is essentially the only formal group of dimension one over \mathcal{O}_ϖ that admits an \mathcal{O}_ϖ -action. (See, e.g., [CF], Chapter VI, §3, for a more detailed discussion of Lubin-Tate formal groups.)

Example 2.8. Suppose that $\lambda_i = \lambda$, for $i = 1, \dots, n$. Then if \mathcal{G}_λ is the Lubin-Tate formal group associated to $\mathcal{O}_\lambda \stackrel{\text{def}}{=} W(\mathbf{F}_{p^\lambda}) \subseteq \mathcal{O}_\varpi$ (since $\varpi = n\lambda$), one checks easily that

$$\mathcal{G}_\Lambda \cong \mathcal{G}_\lambda \otimes_{\mathcal{O}_\lambda}^{\text{gp}} \mathcal{O}_\varpi$$

Indeed, for $j = 0, \dots, \lambda - 1$, let $\epsilon_j \in M_\Lambda$ be the element whose component in $(M_\Lambda)_{i,j'} = A$ (for $i = 1, \dots, n$) is 1 if $j' = j$ and 0 if $j' \neq j$; let $M_\lambda \subseteq M_\Lambda$ be the A -submodule generated by the ϵ_j 's. Then M_λ inherits a Hodge filtration and Frobenius action from M_Λ . Moreover, relative to this Hodge filtration and Frobenius action, M_λ may be identified with the Dieudonné module of G_λ . Finally, the \mathcal{O}_ϖ -action (respectively, \mathcal{O}_λ -action) on \mathcal{G}_Λ (respectively, \mathcal{G}_λ) induces an isomorphism

$$\mathcal{O}_\varpi \otimes_{\mathcal{O}_\lambda} M_\lambda \cong M_\Lambda$$

Thus, since we are using “contravariant Dieudonné modules” and (as \mathcal{O}_ϖ is étale over \mathcal{O}_λ) $\mathcal{O}_\varpi \cong \text{Hom}_{\mathcal{O}_\lambda}(\mathcal{O}_\varpi, \mathcal{O}_\lambda)$, this isomorphism induces an isomorphism $\mathcal{G}_\Lambda \cong \mathcal{G}_\lambda \otimes_{\mathcal{O}_\lambda}^{\text{gp}} \mathcal{O}_\varpi$, as claimed.

Thus, in general, one may think of \mathcal{G}_Λ as some sort of *twisted version* of the product of the \mathcal{G}_{λ_i} 's.

§2.2. The Canonical Uniformizing p-divisible Group

In this subsection, we would like to define a *uniformizing* \mathcal{MF}^∇ -object associated to the Frobenius system. We do this as follows: For $i = 1, \dots, n$, let

$$P_i \stackrel{\text{def}}{=} M_i \oplus \mathcal{O}_S$$

We give P_i a *Hodge filtration* by letting

$$F^0(P_i) \stackrel{\text{def}}{=} P_i; \quad F^1(P_i) \stackrel{\text{def}}{=} F^1(M_i) \oplus 0; \quad F^2(P_i) \stackrel{\text{def}}{=} 0$$

Next, we put a *connection* on P_i as follows: We start with the connection ∇'_{P_i} on P_i which is the *direct sum* of the connection ∇_{M_i} on M_i and the trivial connection on \mathcal{O}_S . Then we *add* to ∇'_{P_i} the $\text{End}(P_i)$ -valued *logarithmic differential* given by the composite

$$P_i \rightarrow M_i \stackrel{\text{def}}{=} \Omega_{S^{\log}/A} \cong (0 \oplus \mathcal{O}_S) \otimes_{\mathcal{O}_S} \Omega_{S^{\log}/A} \subseteq P_i \otimes_{\mathcal{O}_S} \Omega_{S^{\log}/A}$$

where the first morphism is the projection on the first direct summand. The resulting logarithmic connection on P_i will be called ∇_{P_i} . Note that the *Kodaira-Spencer morphism* for $F^1(P_i) \subseteq P_i$ with respect to ∇_{P_i} is the identity map. This property of the Kodaira-Spencer morphism of the Hodge filtration is what will make our \mathcal{MF}^∇ -object *uniformizing*.

Next, we consider *Frobenius*. First, observe that

$$\mathcal{P}_i \stackrel{\text{def}}{=} (P_i, \nabla_{P_i})$$

defines a crystal on $\text{Crys}(S_{\mathbf{F}_p}^{\log}/A)$. Thus, we can form the *renormalized Frobenius pull-back* $\mathbf{F}^*(\mathcal{P}_i)$, i.e., the subsheaf of $\Phi_{S_{\mathbf{F}_p}}^*(\mathcal{P}_i) \otimes_{\mathbf{Z}_p} (\frac{1}{p} \cdot \mathbf{Z}_p)$ of sections whose reductions modulo p lie in $\Phi_{S_{\mathbf{F}_p}}^*(F^1(P_i)_{\mathbf{F}_p}) \otimes (\frac{1}{p})$. Now observe that we have an isomorphism

$$\Phi_{\mathcal{P}_i} : (\Phi_{S_{\mathbf{F}_p}}^*)^{\lambda_i-1} \mathbf{F}^*(\mathcal{P}_i) \cong \mathcal{P}_{i+1}$$

(where we regard \mathcal{P}_{n+1} as \mathcal{P}_1) given as follows: The *left-hand side* $(\Phi_{S_{\mathbf{F}_p}}^*)^{\lambda_i-1} \mathbf{F}^*(\mathcal{P}_i)$ is simply

$$\Phi_i^* P_i = \Phi_i^*(M_i \oplus \mathcal{O}_S)$$

while the *right-hand side* is just

$$P_{i+1} = M_{i+1} \oplus \mathcal{O}_S$$

so by considering the morphism

$$(\Omega_{\Phi_i} \oplus \Phi_i^{-1}) : M_i \oplus \mathcal{O}_S \rightarrow M_{i+1} \oplus \mathcal{O}_S$$

we obtain an isomorphism $\Phi_{\mathcal{P}_i}$ as desired.

Now, for $i = 1, \dots, n$; $j = 0, \dots, \lambda_i - 1$, set

$$\mathcal{P}_{i,j} = (P_{i,j}, \nabla_{P_{i,j}}) \stackrel{\text{def}}{=} \mathcal{P}_i$$

if $j = 0$, and

$$\mathcal{P}_{i,j} = (P_{i,j}, \nabla_{P_{i,j}}) \stackrel{\text{def}}{=} (\Phi_{S_{\mathbf{F}_p}}^*)^{j-1} \mathbf{F}^*(\mathcal{P}_i)$$

if $j > 0$. Next, let us observe that we have *inclusions*

$$\iota_{i,j} : \mathcal{O}_S \hookrightarrow P_{i,j}$$

(via the second direct summand of $P_i = M_i \oplus \mathcal{O}_S$ if $j = 0$; via some Frobenius pull-back of (the horizontal morphism) $\iota_{i,0}$ if $j > 0$) which respect the Hodge filtrations, the connections ∇_{P_i} , and the morphisms Φ_{P_i} . (Here, we give \mathcal{O}_S the Hodge filtration, connection, and Frobenius action of the trivial \mathcal{MF}^∇ -object.) Let

$$\mathcal{P} = (P, \nabla_P)$$

be the inductive limit (or “push-out”) of the following diagram:

$$\begin{array}{ccc} \bigoplus_{i=1}^n \left(\bigoplus_{j=0}^{\lambda_i-1} \mathcal{O}_S \right) & \xrightarrow{\{\iota_{i,j}\}} & \bigoplus_{i=1}^n \left(\bigoplus_{j=0}^{\lambda_i-1} \mathcal{P}_{i,j} \right) \\ \downarrow & & \\ \mathcal{O}_S & & \end{array}$$

where the morphism on the left is the one whose restriction to each direct summand is the identity $\mathcal{O}_S \rightarrow \mathcal{O}_S$. Thus, the \mathcal{O}_S in the bottom left-hand corner of the diagram gives an injection

$$\iota : \mathcal{O}_S \hookrightarrow P$$

Moreover, P inherits a *Hodge filtration* from those of the $P_{i,j}$. Finally, the Φ_{P_i} clearly define a *Frobenius action* Φ_P on P so that we get an \mathcal{MF}^∇ -object

$$(P, \nabla_P, F^\bullet(P), \Phi_P)$$

Moreover, the morphism $\iota : \mathcal{O}_S \rightarrow P$ defines a morphism of the trivial \mathcal{MF}^∇ -object into the above \mathcal{MF}^∇ -object such that the cokernel of this morphism is the \mathcal{MF}^∇ -object

$$(M, \nabla_M, F^\bullet(M), \Phi_M)$$

of the discussion preceding Definition 2.3. Now because of the way we defined the Hodge filtration, $(P, \nabla_P, F^\bullet(P), \Phi_P)$ defines a *log p-divisible group* G_Φ on S^{\log} (cf. [Falt1], §7, Theorem 7.1 – at least for the “non-logarithmic case”; for more on the logarithmic case, cf. [Mzk1], Chapter IV, §2). Moreover, this log p -divisible group fits into an exact sequence

$$0 \rightarrow G_{\Omega_\Phi} \rightarrow G_\Phi \rightarrow \mathbf{Q}_p/\mathbf{Z}_p \rightarrow 0$$

Let

$$(T'_m)^{\log} \rightarrow S^{\log}$$

be the *finite flat morphism which parametrizes splittings of the above exact sequence over $(\mathbf{Q}_p/\mathbf{Z}_p)[p^m]$* . Since multiplication by p on the formal group \mathcal{G}_Λ is of degree p^ϖ , it follows (cf. Proposition 2.4) that $(T'_m)^{\log} \rightarrow S^{\log}$ is of degree $p^{m \cdot \varpi \cdot d}$. Let

$$\mathcal{U}'_m$$

be the finite $\Pi_{S^{\log}}$ -set corresponding to $((T'_m)^{\log})_K \rightarrow S_K^{\log}$. Let

$$\mathcal{U}' \stackrel{\text{def}}{=} \varprojlim \mathcal{U}'_m$$

where the inverse limit is over integers $m \geq 1$.

Now we make the following

Definition 2.9. We shall call G_Φ the *canonical uniformizing (log) p -divisible group associated to the given Frobenius system*.

To some extent,

When working with systems involving more than one Frobenius lifting, it is more natural to work on a product of copies of S .

To this end, we introduce the following objects: Let

$$S_{\text{PD}}^{\log}$$

be the p -adic completion of the PD-envelope of $\prod S^{\log}$ (where the product is the product of n copies of S^{\log} , numbered $1, \dots, n$) with respect to the diagonal $S_{\mathbf{F}_p}^{\log} \hookrightarrow \prod S^{\log}$ modulo p . Let

$$S_{\text{FM}}^{\log}$$

be the formal log scheme defined by completing $\prod S^{\log}$ at the diagonal. Thus, we have diagonal embeddings

$$S^{\log} \hookrightarrow S_{\text{PD}}^{\log}; \quad S^{\log} \hookrightarrow S_{\text{FM}}^{\log}$$

as well as a morphism

$$S_{\text{PD}}^{\log} \rightarrow S_{\text{FM}}^{\log}$$

compatible with the diagonal embeddings. By letting Φ_σ act on the i^{th} factor (where $\sigma(1) = i$), we see that the Φ_σ 's define a morphism

$$\Phi_{S_{\text{FM}}}^{\log} : S_{\text{FM}}^{\log} \rightarrow S_{\text{FM}}^{\log}$$

whose reduction modulo p is the ϖ^{th} power of Frobenius. Note that $\Phi_{S_{\text{FM}}}^{\log}$ also extends to a morphism

$$\Phi_{S_{\text{PD}}}^{\log} : S_{\text{PD}}^{\log} \rightarrow S_{\text{PD}}^{\log}$$

Ideally, we would like to extend the uniformizing group G_Φ to a p -divisible group on S_{FM}^{\log} in a natural way. Unfortunately, in general, if $n \geq 2$, such an extension simply does not exist (as we shall see later – cf. §2.6). Thus, we content ourselves here with constructing an extension of G_Φ to S_{PD}^{\log} .

First, let us note that in fact,

We can construct something very close to an \mathcal{MF}^∇ -object on S_{FM}^{\log} whose restriction to the diagonal $S^{\log} \hookrightarrow S_{\text{FM}}^{\log}$ gives rise to G_Φ .

Indeed, let

$$\pi_i^{\log} : S_{\text{FM}}^{\log} \rightarrow S^{\log}$$

be the projection to the i^{th} factor. Then we let

$$\mathcal{P}_{S_{\text{FM}}} = (P_{S_{\text{FM}}}, \nabla_{P_{S_{\text{FM}}}})$$

be the inductive limit of the following diagram:

$$\begin{array}{ccc} \bigoplus_{i=1}^n \left(\bigoplus_{j=0}^{\lambda_i-1} \mathcal{O}_S \right) & \xrightarrow{\{\iota_{i,j}\}} & \bigoplus_{i=1}^n \left(\bigoplus_{j=0}^{\lambda_i-1} \pi_i^* \mathcal{P}_{i,j} \right) \\ \downarrow & & \\ \mathcal{O}_S & & \end{array}$$

where the arrows are defined similarly to the way they were defined in the definition of \mathcal{P} . Note that by pulling back the Hodge filtrations on the $\mathcal{P}_{i,j}$ via π_i , we get a Hodge filtration on $\mathcal{P}_{S_{\text{FM}}}$.

Now the problem is to define a Frobenius action on $\mathcal{P}_{S_{\text{FM}}}$. Unfortunately, this is not possible in general. However, it is possible to define the “ ϖ^{th} power of the Frobenius action,” i.e., a horizontal morphism

$$(\Phi_{S_{\text{FM}}})_P : \Phi_{S_{\text{FM}}}^* P_{S_{\text{FM}}} \rightarrow P_{S_{\text{FM}}}$$

(which becomes an isomorphism after one adjusts the integral structure on $\Phi_{S_{\text{FM}}}^* P_{S_{\text{FM}}}$ in some appropriate fashion). Indeed, we define $(\Phi_{S_{\text{FM}}})_P$ on the portion of $P_{S_{\text{FM}}}$ arising from $\pi_i^* P_{i,0}$ by letting it be the (pull-back by π_i of the) composite of $\Phi_{\mathcal{P}_{\sigma(n)}}, \dots, \Phi_{\mathcal{P}_{\sigma(1)}}$ (where $\sigma(1) = i$). We then extend $(\Phi_{S_{\text{FM}}})_P$ to the $\pi_i^* P_{i,j}$ (for $j > 0$) by pulling back $(\Phi_{S_{\text{FM}}})_P$ by the same combination of renormalized and naive Frobenius pull-backs that were used to construct $\mathcal{P}_{i,j}$ out of $\mathcal{P}_{i,0}$. Note that with this definition, $(\Phi_{S_{\text{FM}}})_P$ is compatible with the connection and Hodge filtration on $P_{S_{\text{FM}}}$, and, moreover, with the Frobenius action on the restriction of $P_{S_{\text{FM}}}$ to the diagonal $S^{\log} \subseteq S_{\text{FM}}^{\log}$. Thus, over S_{FM}^{\log} , we have constructed the following data

$$(P_{S_{\text{FM}}}, \nabla_{P_{S_{\text{FM}}}}, F^*(P_{S_{\text{FM}}}), (\Phi_{S_{\text{FM}}})_P)$$

Now let us consider what happens when one pulls the above data back to S_{PD}^{\log} from S_{FM}^{\log} . Note that since S_{PD}^{\log} is a PD-thickening of $S_{\mathbf{F}_p}^{\log}$ (over A), it follows that the restriction

$$\mathcal{P}_{S_{\text{PD}}} \stackrel{\text{def}}{=} \mathcal{P}_{S_{\text{FM}}} |_{S_{\text{PD}}^{\log}}$$

is completely determined by the restriction of $\mathcal{P}_{S_{\text{FM}}}$ to the diagonal $S^{\log} \subseteq S_{\text{FM}}^{\log}$, i.e., by the crystal $\mathcal{P} = (P, \nabla_P)$ on S^{\log} . Thus, the existence of a Frobenius action on \mathcal{P} means that there is a Frobenius action

$$\Phi_{P_{S_{\text{PD}}}}$$

on $\mathcal{P}_{S_{\text{PD}}}$. Moreover, it is easy to see that (if we forget about the factors of $\frac{1}{p}$ used to modify the integral structures) then $\{\Phi_{P_{S_{\text{PD}}}}\}^{\varpi}$ may be identified with $(\Phi_{S_{\text{FM}}})_P |_{S_{\text{PD}}}$. In summary, over S_{PD}^{\log} , we have the following data

$$(P_{S_{\text{PD}}}, \nabla_{P_{S_{\text{PD}}}}, F^*(P_{S_{\text{PD}}}), \Phi_{P_{S_{\text{PD}}}})$$

Now we would like to claim that this data gives rise to a (log) p -divisible group on S_{PD}^{\log} . Unfortunately, we cannot apply [Falt1], §7, Theorem 7.1, directly, because the base S_{PD}^{\log} is not of the type considered in [Falt1]. It is, however, a PD-thickening of S^{\log} , which is of the type considered in [Falt1]. Thus, we can apply [Mess], Chapter V, Theorem 1.6, which states that:

The deformations of a p -divisible group to a PD-thickening of a given base are categorically equivalent to deformations of the Hodge filtration of the Dieudonné crystal of the p -divisible group.

(Note that in [Mess], Chapter V, it is assumed that the p -divisible groups in question all lift Zariski locally on the base. In the present context, this assumption is satisfied because, for instance, $\pi_i^* G_\Phi$ (for any $i = 1, \dots, n$) restricts to G_Φ on the diagonal $S^{\log} \subseteq S_{\text{PD}}^{\log}$.) In particular, it follows that the data just defined forms the filtered Dieudonné module associated to a (log) p -divisible group

$$G_{\Phi_{\text{PD}}}$$

on S_{PD} . Moreover, we have (by construction):

$$G_{\Phi_{\text{PD}}} |_{S^{\log}} = G_\Phi$$

(where $S^{\log} \subseteq S_{\text{PD}}^{\log}$ by the diagonal embedding).

Definition 2.10. We shall call $G_{\Phi_{\text{PD}}}$ the canonical multi-uniformizing (log) p -divisible group associated to the given Frobenius system.

Next, note that since S_{PD} is a (topologically) nilpotent thickening of the diagonal S (hence has the same étale site as S), and the differential p -divisible group G_{Ω_Φ} (of Definition 2.3) is obtained by tensoring (in the sense of \otimes^{gp}) an étale p -divisible group with a p -divisible group pulled back from $\text{Spec}(A)$ (cf. Proposition 2.4), it follows that G_{Ω_Φ} extends immediately to S_{PD}^{\log} . We denote the resulting group on S_{PD}^{\log} by $G_{\Omega_\Phi} |_{S_{\text{PD}}^{\log}}$. Thus, we have an exact sequence

$$0 \rightarrow G_{\Omega_\Phi} |_{S_{\text{PD}}^{\log}} \rightarrow G_{\Phi_{\text{PD}}} \rightarrow \mathbf{Q}_p / \mathbf{Z}_p \rightarrow 0$$

We would like to relate this extension to the Frobenius lifting $\Phi_{S_{\text{PD}}}^{\log} : S_{\text{PD}}^{\log} \rightarrow S_{\text{PD}}^{\log}$. First note that since $G_{\Omega_\Phi} |_{S_{\text{PD}}^{\log}}$ essentially comes from $\text{Spec}(A)$, the pull-back of $G_{\Omega_\Phi} |_{S_{\text{PD}}^{\log}}$ via $\Phi_{S_{\text{PD}}}$ is naturally isomorphic to itself. Moreover, since the morphism $(\Phi_{S_{\text{FM}}})_P |_{S_{\text{PD}}}$ is compatible with connections, Hodge filtrations, and the Frobenius action (cf. the discussion following the definition of $(\Phi_{S_{\text{FM}}})_P$), it follows from the fully-faithfulness of the Dieudonné crystal functor on p -divisible groups (cf., e.g., [Falt1], §7, Theorem 7.1; [Mess], Chapter V, Theorem 1.6) that we get a commutative diagram of extensions as follows:

$$\begin{array}{ccccccc}
0 & \longrightarrow & G_{\Omega_{\Phi}}|_{S_{\text{PD}}^{\log}} & \longrightarrow & G_{\Phi_{\text{PD}}} & \longrightarrow & \mathbf{Q}_p/\mathbf{Z}_p \longrightarrow 0 \\
& & \downarrow p^n & & \downarrow & & \downarrow \text{id} \\
0 & \longrightarrow & G_{\Omega_{\Phi}}|_{S_{\text{PD}}^{\log}} & \longrightarrow & \Phi_{S_{\text{PD}}}^*(G_{\Phi_{\text{PD}}}) & \longrightarrow & \mathbf{Q}_p/\mathbf{Z}_p \longrightarrow 0
\end{array}$$

In other words, (by considering the corresponding commutative diagram obtained by replacing each of the p -divisible groups in the above diagram by its corresponding “ $(-)[p^n]$,” i.e., the kernel of multiplication by p^n) we see that *the exact sequence*

$$0 \rightarrow G_{\Omega_{\Phi}}[p^n]|_{S_{\text{PD}}^{\log}} \rightarrow G_{\Phi_{\text{PD}}}[p^n] \rightarrow \mathbf{Q}_p/\mathbf{Z}_p[p^n] \rightarrow 0$$

splits after pull-back by $\Phi_{S_{\text{PD}}}$.

Next, we would like to begin to *relate the uniformizing p -divisible groups defined above to the original system of Frobenius liftings*. More precisely, we would like to show that there is a natural isomorphism (of $\Pi_{S^{\log}}$ -sets) between \mathcal{U}_m (cf. Definition 1.4) and \mathcal{U}'_{nm} (cf. the discussion preceding Definition 2.9). We begin by *constructing a morphism*

$$\mathcal{U}_m \rightarrow \mathcal{U}'_{nm}$$

(of $\Pi_{S^{\log}}$ -sets). To construct such a morphism, it suffices to construct a morphism

$$\beta_m : (T_m^{\log})_K \rightarrow (T'_{nm})_K^{\log}$$

(over S_K^{\log}) where $(T'_{nm})^{\log}$ is as in the discussion preceding Definition 2.9, and T_m^{\log} is defined by the following cartesian diagram:

$$\begin{array}{ccc}
T_m^{\log} & \longrightarrow & S_{\text{FM}}^{\log} \\
\downarrow & & \downarrow (\Phi_{S_{\text{FM}}}^{\log})^m \\
S^{\log} & \longrightarrow & S_{\text{FM}}^{\log}
\end{array}$$

(Note that it follows from the definitions that \mathcal{U}_m is precisely the $\Pi_{S^{\log}}$ -set corresponding to $T_m^{\log} \rightarrow S^{\log}$.) To construct β_m , it thus suffices to construct a *splitting of the pull-back to $(T_m^{\log})_K$ (from S^{\log}) of the exact sequence*

$$0 \rightarrow G_{\Omega_{\Phi}}[p^{nm}] \rightarrow G_{\Phi}[p^{nm}] \rightarrow \mathbf{Q}_p/\mathbf{Z}_p[p^{nm}] \rightarrow 0$$

We shall construct such a splitting of group schemes by *constructing a splitting of the corresponding Dieudonné crystals*. To do this, we would first like to consider the pull-back morphism

$$\{(\Phi_{S_{\mathbf{F}_p}})^{m \cdot \varpi}\}^{-1} : \text{Crys}(S_{\mathbf{F}_p}^{\log}/A) \rightarrow \text{Crys}(S_{\mathbf{F}_p}^{\log}/A)$$

between crystalline sites induced by the $(m \cdot \varpi)^{\text{th}}$ power of the Frobenius morphism $\Phi_{S_{\mathbf{F}_p}}$ on $S_{\mathbf{F}_p}$. Then for $i = 1, \dots, n$, we have a sub-crystal

$$\mathcal{F}_{i,0} \subseteq \{(\Phi_{S_{\mathbf{F}_p}})^{m \cdot \varpi}\}^*(\mathcal{P}_{i,0})_{\mathbf{Z}/p^{nm}\mathbf{Z}}$$

induced by the \mathcal{O}_S -submodule

$$(\Phi_{\sigma}^m)^* F^1(P_{i,0})_{\mathbf{Z}/p^{nm}\mathbf{Z}} \subseteq (\Phi_{\sigma}^m)^*(P_{i,0})_{\mathbf{Z}/p^{nm}\mathbf{Z}}$$

(where $\sigma(1) = i$). Note that this \mathcal{O}_S -submodule defines a *sub-crystal* since the Kodaira-Spencer morphism of this $F^1(-)$ becomes $\equiv 0$ modulo p^{nm} after pull-back by Φ_{σ}^m . Now let us denote by

$$\xi_i^{-1} : \text{Crys}(S_{\mathbf{F}_p}^{\log}/A) \rightarrow \text{Crys}((T_m^{\log})_{\mathbf{F}_p}/A)$$

the pull-back morphism on crystalline sites induces by the restriction to $T_m^{\log} \subseteq S_{\text{FM}}$ of the projection morphism $\pi_i : S_{\text{FM}} \rightarrow S$. Thus, by pulling back $\mathcal{F}_{i,0}$ by ξ_i^{-1} , we obtain a sub-crystal

$$\mathcal{G}_{i,0} \subseteq (\mathcal{P}_{i,0})_{\mathbf{Z}/p^{nm}\mathbf{Z}}|_{T_m^{\log}}$$

(where we use the notation “ $|_{T_m^{\log}}$ ” since the composite of ξ_i^{-1} with $\{(\Phi_{S_{\mathbf{F}_p}})^{m \cdot \varpi}\}^{-1}$ is clearly independent of i). Note that for $j = 1, \dots, \lambda_i - 1$, by forming various renormalized (if $j = 1$) or naive (if $j > 1$) Frobenius pull-backs of $\mathcal{G}_{i,0} \subseteq (\mathcal{P}_{i,0})_{\mathbf{Z}/p^{nm}\mathbf{Z}}|_{T_m^{\log}}$, we obtain $\mathbf{Z}/p^{nm}\mathbf{Z}$ -flat sub-crystals

$$\mathcal{G}_{i,j} \subseteq (\mathcal{P}_{i,j})_{\mathbf{Z}/p^{nm}\mathbf{Z}}|_{T_m^{\log}}$$

Finally, I *claim* that if we form the renormalized (if $\lambda_i = 1$) or naive (if $\lambda_i > 1$) Frobenius pull-back of $\mathcal{G}_{i,j} \subseteq (\mathcal{P}_{i,j})_{\mathbf{Z}/p^{nm}\mathbf{Z}}|_{T_m^{\log}}$, we obtain

$$\mathcal{G}_{i+1,0} \subseteq (\mathcal{P}_{i+1,0})_{\mathbf{Z}/p^{nm}\mathbf{Z}}|_{T_m^{\log}}$$

Indeed, to see this, it suffices to compute the relevant Frobenius pull-back of $\mathcal{F}_{i,0}$ using the Frobenius lifting Φ_i . But then the claim follows from the fact that

$$\Phi_\sigma \circ \Phi_i = (\Phi_i \circ \Phi_{i+1} \dots \circ \Phi_{i-1}) \circ \Phi_i = \Phi_i \circ (\Phi_{i+1} \circ \dots \circ \Phi_i) = \Phi_i \circ \Phi_{\sigma'}$$

(where $\sigma'(1) = i + 1$), together with the fact that the morphism $\Phi_i^* P_{i,0} \rightarrow P_{i+1,0}$ (used to define the Frobenius action on \mathcal{P}) *preserves the Hodge filtration*.

Let us now pause to review what has been accomplished. We have defined sub-crystals (for $i = 1, \dots, n$; $j = 0, \dots, \lambda_i - 1$)

$$\mathcal{G}_{i,j} \subseteq (\mathcal{P}_{i,j})_{\mathbf{Z}/p^{nm}\mathbf{Z}}|_{T_m^{\log}}$$

in such a way that the image of

$$\mathcal{G} \stackrel{\text{def}}{=} \bigoplus_{i,j} \mathcal{G}_{i,j}$$

in $\mathcal{P}_{\mathbf{Z}/p^{nm}\mathbf{Z}}|_{T_m^{\log}}$ defines a sub-crystal

$$\mathcal{G} \subseteq \mathcal{P}_{\mathbf{Z}/p^{nm}\mathbf{Z}}|_{T_m^{\log}}$$

which is invariant under the Frobenius action on $\mathcal{P}_{\mathbf{Z}/p^{nm}\mathbf{Z}}|_{T_m^{\log}}$ and, moreover, *maps isomorphically onto $\mathcal{M}_{\mathbf{Z}/p^{nm}\mathbf{Z}}|_{T_m^{\log}}$* (via the natural projection $\mathcal{P} \rightarrow \mathcal{M}$). If we then identify \mathcal{G} with $\mathcal{M}_{\mathbf{Z}/p^{nm}\mathbf{Z}}|_{T_m^{\log}}$ (via this isomorphism), it follows from our definition of the $\mathcal{F}_{i,0}$ (*involving various $F^1(-)$'s!*) that the original Hodge filtration on $\mathcal{M}_{\mathbf{Z}/p^{nm}\mathbf{Z}}|_{T_m^{\log}}$ coincides with the Hodge filtration induced by the inclusion $\mathcal{M}_{\mathbf{Z}/p^{nm}\mathbf{Z}}|_{T_m^{\log}} \subseteq \mathcal{P}_{\mathbf{Z}/p^{nm}\mathbf{Z}}|_{T_m^{\log}}$. In short, we have shown that the pull-back of $\mathcal{P}_{\mathbf{Z}/p^{nm}\mathbf{Z}}$ to T_m^{\log} *admits a decomposition*

$$\mathcal{P}_{\mathbf{Z}/p^{nm}\mathbf{Z}}|_{T_m^{\log}} = \mathcal{O}_{(T_m^{\log})_{\mathbf{Z}/p^{nm}\mathbf{Z}}} \oplus (\mathcal{M}_{\mathbf{Z}/p^{nm}\mathbf{Z}}|_{T_m^{\log}})$$

that respects the Frobenius actions and Hodge filtrations on both sides.

On the other hand, it follows from the definition of the “Galois representation associated to an \mathcal{MF}^∇ -object” (cf. [Falt1], §2, e)) that such a direct sum decomposition implies that the pull-back to $(T_m^{\log})_K$ of the exact sequence

$$0 \rightarrow G_{\Omega_\Phi}[p^{nm}] \rightarrow G_\Phi[p^{nm}] \rightarrow \mathbf{Q}_p/\mathbf{Z}_p[p^{nm}] \rightarrow 0$$

splits, as desired. This completes the definition of the morphism

$$\beta_m : (T_m^{\log})_K \rightarrow (T'_{nm})_K^{\log}$$

Of course, we would like to show that β_m is an *isomorphism*. To this end, note first of all that both $(T_m^{\log})_K$ and $(T'_{nm})_K^{\log}$ are *finite étale* over S_K^{\log} , of degree $p^{nmd \cdot \varpi}$. Thus, if we knew, for instance, that $(T'_{nm})_K^{\log}$ is *connected*, then it would follow immediately from elementary properties of finite étale morphisms that β_m is an isomorphism.

In fact, for an arbitrary system of Frobenius liftings $\Phi_1^{\log}, \dots, \Phi_n^{\log}$, it is not always the case that $(T'_{nm})_K^{\log}$ is connected. In the following few paragraphs, we would like to discuss a certain *sufficient condition for the bijectivity of β_m* (related to the connectedness of $(T'_{nm})_K^{\log}$). Then, we shall quote Lemma 2.11 below, which states that it is always possible to *deform* the given Frobenius liftings $\Phi_1^{\log}, \dots, \Phi_n^{\log}$ (over an irreducible base space) to a system of Frobenius liftings for which this sufficient condition is satisfied. By elementary properties of finite étale morphisms, one knows that the issue of whether or not β_m is bijective is *locally constant* on the base space. Thus, we conclude that β_m is always bijective (Theorem 2.12 below), as desired.

For simplicity, let us assume (without loss of generality – as one sees by replacing S by an appropriate finite étale cover of S) in the following discussion that the étale local system $\Omega_{\Phi_\sigma}^{\text{et}} \otimes \mathbf{Z}/p^{nm}\mathbf{Z}$ is *trivial* (for some, or equivalently, for all shifting permutations σ). Write

$$\Delta_{S^{\log}} \stackrel{\text{def}}{=} \text{Ker}(\Pi_{S^{\log}} \rightarrow \text{Gal}(\overline{K}/K)) \subseteq \Pi_{S^{\log}}$$

for the “geometric fundamental group” of S_K^{\log} . Now β_m defines a morphism of $\Pi_{S^{\log}}$ -sets

$$\mathcal{U}_m \rightarrow \mathcal{U}'_{nm} \quad \bullet$$

Let $\mathcal{V} \subseteq \mathcal{U}'_{nm}$ be the *image* of this morphism. Thus, \mathcal{V} is stable under the action of $\Pi_{S^{\log}}$. Let $v \in \mathcal{V}$. Let \mathcal{A} be the $\Pi_{S^{\log}}$ -module corresponding to $G_{\Omega_\Phi}[p^{nm}]_K$. Note that because of the assumption on $\Omega_{\Phi_\sigma}^{\text{et}} \otimes \mathbf{Z}/p^{nm}\mathbf{Z}$ (and Proposition 2.4), it follows that $\Delta_{S^{\log}} \subseteq \Pi_{S^{\log}}$ *acts trivially* on \mathcal{A} . Thus, since \mathcal{U}'_{nm} is a torsor (i.e., a principal homogeneous space) under \mathcal{A} , we obtain a *twisted homomorphism*

$$H : \Pi_{S^{\log}} \rightarrow \mathcal{A}$$

given by $\pi \cdot v = v + H(\pi)$ (for $\pi \in \Pi_{S^{\log}}$) and satisfying $H(\pi \cdot \pi') = H(\pi) + \pi \cdot H(\pi')$. Let

$$\mathcal{B} \stackrel{\text{def}}{=} H(\Delta_{S^{\log}}) \subseteq \mathcal{A}$$

Since $\Delta_{S^{\log}}$ acts trivially on \mathcal{A} , the restriction of H to $\Delta_{S^{\log}}$ is a homomorphism, so it follows that \mathcal{B} is a \mathbf{Z}_p -submodule of \mathcal{A} . Moreover, if $\delta \in \Delta_{S^{\log}}$, $\pi \in \Pi_{S^{\log}}$, we have:

$$\begin{aligned} H(\pi \cdot \delta \cdot \pi^{-1}) &= H(\pi) + \pi \cdot H(\delta \cdot \pi^{-1}) \\ &= H(\pi) + \pi \cdot \{H(\delta) + \delta \cdot H(\pi^{-1})\} \\ &= H(\pi) + \pi \cdot H(\delta) + \pi \cdot H(\pi^{-1}) \\ &= \pi \cdot H(\delta) + \{H(\pi) + \pi \cdot H(\pi^{-1})\} \\ &= \pi \cdot H(\delta) \end{aligned}$$

Thus, \mathcal{B} is a $\Pi_{S^{\log}}$ -submodule of \mathcal{A} . Moreover, by the definition of \mathcal{B} , it follows that if we push the \mathcal{A} -torsor \mathcal{U}'_{nm} forward by the “change of structure group” $\mathcal{A} \rightarrow \mathcal{A}/\mathcal{B}$, the resulting \mathcal{A}/\mathcal{B} -torsor admits a $(\Pi_{S^{\log}}\text{-invariant})$ trivialization.

Now by basic properties of crystalline representations (cf., e.g., [Falt1], §2, Theorem 2.6), it thus follows that the quotient $\mathcal{A} \otimes \mathbf{F}_p \rightarrow (\mathcal{A}/\mathcal{B}) \otimes \mathbf{F}_p$ corresponds to some \mathcal{MF}^{∇} -subobject

$$(N, \nabla_N, F^{\cdot}(N), \Phi_N) \subseteq (M, \nabla_M, F^{\cdot}(M), \Phi_M)_{\mathbf{F}_p}$$

(annihilated by p). Moreover, since the finite flat group scheme $G_{\Omega_{\Phi}}[p]$ is connected, it follows that if $N \neq 0$, then $F^1(N) \neq 0$. On the other hand, what we did in the preceding paragraph implies (by basic properties of crystalline representations – cf., e.g., [Falt1], §2, Theorem 2.6) that if we pull-back the surjection

$$(P, \nabla_P, F^{\cdot}(P), \Phi_P)_{\mathbf{F}_p} \rightarrow (M, \nabla_M, F^{\cdot}(M), \Phi_M)_{\mathbf{F}_p}$$

by

$$(N, \nabla_N, F^{\cdot}(N), \Phi_N) \subseteq (M, \nabla_M, F^{\cdot}(M), \Phi_M)_{\mathbf{F}_p}$$

the resulting \mathcal{MF}^{∇} -object decomposes as the direct sum of “ $\mathcal{O}_{S_{\mathbf{F}_p}}$ ” (i.e., the trivial rank one \mathcal{MF}^{∇} -object annihilated by p) and $(N, \nabla_N, F^{\cdot}(N), \Phi_N)$. In particular, it follows that the Kodaira-Spencer morphism

$$\kappa : F^1(M)_{\mathbf{F}_p} = \bigoplus_{i=1}^n \Omega_{S_{\mathbf{F}_p}}^{\log} \rightarrow \Omega_{S_{\mathbf{F}_p}}^{\log}$$

(for the Hodge filtration of $\mathcal{P}_{\mathbf{F}_p}$) becomes zero upon restriction to $F^1(N) \subseteq F^1(M)_{\mathbf{F}_p}$. On the other hand, since $F^1(N) = N \cap F^1(M)_{\mathbf{F}_p}$ is stabilized

by the ϖ^{th} power of the Frobenius action on \mathcal{M} , it follows that $F^1(N) \subseteq F^1(M)_{\mathbf{F}_p}$ is generated over $\mathcal{O}_{S_{\mathbf{F}_p}}$ by some \mathbf{F}_p^{ϖ} -subspace of

$$\bigoplus_{\sigma} (\Omega_{\Phi_{\sigma}}^{\text{et}} \otimes \mathbf{F}_p)$$

(where the direct sum is over all the shifting permutations σ). Let us denote the restriction of κ to this direct sum by

$$\kappa^{\text{et}} : \bigoplus_{\sigma} (\Omega_{\Phi_{\sigma}}^{\text{et}} \otimes \mathbf{F}_p) \rightarrow \Omega_{S_{\mathbf{F}_p}}^{\log}$$

Thus, the injectivity of κ^{et} would imply that $F^1(N)$, hence also N , is 0.

Let us summarize what we have done so far. We have shown that the injectivity of κ^{et} implies that $N = 0$, hence that $(\mathcal{A}/\mathcal{B}) \otimes \mathbf{F}_p = 0$. But this implies that $\mathcal{B} = \mathcal{A}$, hence that β_m is surjective. Since β_m is a morphism between finite étale morphisms of the same degree, this implies that β_m is an isomorphism, as desired. In other words,

The injectivity of κ^{et} implies that β_m is an isomorphism.

On the other hand, we have the following:

Lemma 2.11. *(After completing S at some point of $S(k)$) there exists a deformation of the given system of Frobenius liftings $\Phi_1^{\log}, \dots, \Phi_n^{\log}$ (over an irreducible base space) to a new system of binary-ordinary Frobenius liftings for which κ^{et} is injective.*

Since the proof of Lemma 2.11 is rather technical, we shall save it for a later subsection (§2.7). At any rate, once one admits Lemma 2.11, we obtain that any binary-ordinary Frobenius system may be deformed (over an irreducible base space) to a system for which β_m is bijective. Since (as noted above) the issue of whether or not β_m is bijective is locally constant on the base, this implies that β_m is always bijective. That is to say, we have proven the following result, relating the original system of Frobenius liftings to its canonical multi-uniformizing p -divisible group:

Theorem 2.12. *For every integer $m \geq 1$, we have natural isomorphisms of $\Pi_{S^{\log}}$ -sets:*

$$\mathcal{U}_m \cong \mathcal{U}'_{nm}; \quad \mathcal{U} \cong \mathcal{U}'$$

In other words, the set-theoretic canonical Galois representation of Definition 1.4 can be naturally recovered from the canonical uniformizing p -divisible group G_Φ .

Remark. In an earlier version of this work ([Mzk5,6]), it was claimed that the isomorphism $(T_m^{\log})_K \rightarrow (T'_{nm})_K^{\log}$ in fact arises from an isomorphism $T_m^{\log} \rightarrow (T'_{nm})^{\log}$. Upon closer inspection, however, the author discovered a gap in this proof. It was precisely in order to fill this gap that the technique of “deforming to a Frobenius system for which κ^{et} is injective” was introduced.

§2.3. Multi-Uniformization by the Group \mathcal{G}_Λ

Let us assume just in this subsection that k is *algebraically closed*. Let $z \in S(k)$ be a k -valued point of S . Let S_z^{\log} be the completion of S^{\log} at z . Thus, S_z^{\log} is Spf of a complete local ring R_z which is noncanonically isomorphic to

$$A[[t_1, \dots, t_d]]$$

with the restriction of the divisor D (that defines the log structure of S) defined by $t_1 \cdot t_2 \cdot \dots \cdot t_a$ (where a may be zero). We shall also write

$$(S_{\text{FM}})_z \quad \text{and} \quad (S_{\text{PD}})_z$$

for the completions of S_{FM} and S_{PD} , respectively, at z (regarded as a point of S_{FM} or S_{PD} via the diagonal embedding). Let $U = S - D$. Let us denote by $\mathcal{G}_\Lambda^{\log}((S_{\text{PD}})_z)$ the following:

(1) If $z \in U$, then

$$\mathcal{G}_\Lambda^{\log}((S_{\text{PD}})_z) \stackrel{\text{def}}{=} \mathcal{G}_\Lambda((S_{\text{PD}})_z)$$

(2) If $z \in D$, then

$$\mathcal{G}_\Lambda^{\log}((S_{\text{PD}})_z) \stackrel{\text{def}}{=} (\mathbf{G}_m((S_{\text{PD}})_z \times_S U))^\wedge \otimes_{\mathbf{Z}_p} \mathcal{O}_\varpi$$

(where the “ \wedge ” denotes p -adic completion).

Note that the definition of (2) is consistent with the definition of (1) when all the $\lambda_i = 1$ (cf. Example 2.8). Since k is *algebraically closed*, if we restrict any of the

$$\Omega_{\Phi_\sigma}^{\text{et}}$$

to S_z^{\log} , we obtain the *trivial local system*. Thus, in the following discussion, we shall (by abuse of notation) denote by $\Omega_{\Phi_\sigma}^{\text{et}}$ the free \mathcal{O}_ϖ -module (of rank $d = \dim_A(S)$) corresponding to this trivial local system.

Let $\omega \in \Omega_{\Phi_{\text{id}}}^{\text{et}}$. Let us denote by

$$\omega_\sigma \in \Omega_{\Phi_\sigma}^{\text{et}}$$

the differential corresponding to ω under the canonical bijection of $\Omega_{\Phi_{\text{id}}}^{\text{et}}$ with $\Omega_{\Phi_\sigma}^{\text{et}}$ defined by the Ω_{Φ_i} . Thus, $\omega_{\text{id}} = \omega$. Now, ω defines a surjection

$$\omega(-) : \Theta_{\Phi_{\text{id}}}^{\text{et}} \rightarrow \mathcal{O}_\varpi$$

If we restrict our canonical extension of (log) p -divisible groups

$$0 \rightarrow G_{\Omega_\Phi} |_{S_{\text{PD}}^{\log}} \rightarrow G_{\Phi_{\text{PD}}} \rightarrow \mathbf{Q}_p/\mathbf{Z}_p \rightarrow 0$$

to $(S_{\text{PD}}^{\log})_z$, and use Proposition 2.4 to apply $\omega(-)$, we obtain a (log) extension

$$0 \rightarrow G_\Lambda \rightarrow G_{\omega,z} \rightarrow \mathbf{Q}_p/\mathbf{Z}_p \rightarrow 0$$

of $\mathbf{Q}_p/\mathbf{Z}_p$ by G_Λ over $(S_{\text{PD}}^{\log})_z$. By the generalization of Kummer theory given in Proposition 2.6, we thus obtain a “logarithmic unit”

$$q_{\omega,z} \in \mathcal{G}_\Lambda^{\log}((S_{\text{PD}})_z)$$

In other words, we have a “morphism”

$$q_{\omega,z} : (S_{\text{PD}}^{\log})_z \rightarrow \mathcal{G}_\Lambda^{\log}$$

(where in the logarithmic case, we must interpret the notion of “morphism” properly – we leave this to the reader) i.e., a \mathcal{G}_Λ -coordinate associated to ω . Now recall the *canonical differentials* $\omega_{\Lambda,\sigma}$ on \mathcal{G}_Λ (cf. (4) of the discussion following Proposition 2.4). Then the natural *question of compatibility* arises: Is

$$q_{\omega,z}^*(\omega_{\Lambda,\sigma}) \stackrel{?}{=} \pi_i^* \omega_\sigma$$

(where $\sigma(1) = i$)?

To answer this question, first we focus our attention on the case where $z \in U$. Thus, suppose (just for the rest of this paragraph) that $z \in U$. Let

$$\alpha \in (S_{\text{PD}})_z(A)$$

Let

$$P_{\omega,\alpha}$$

be the portion of the Dieudonné crystal of $G_{\omega,z}|_{\alpha}$ generated by the “ $P_{i,0}$ ” (for $i = 1, \dots, n$) (cf. the construction of P preceding Definition 2.9). Thus, $P_{\omega,\alpha}$ is a free A -module of rank $1 + n$ equipped with a filtration $F^1(P_{\omega,\alpha}) \subseteq P_{\omega,\alpha}$, and a “ ϖ^{th} power of Frobenius” action

$$\Phi_{\omega,\alpha} : P_{\omega,\alpha}^{F^{\varpi}} \rightarrow P_{\omega,\alpha}$$

This Frobenius action has a unique subspace

$$P_F \subseteq P_{\omega,\alpha} \quad (\text{respectively, } P_V \subseteq P_{\omega,\alpha})$$

on which Frobenius acts with slope zero (respectively, slope > 0). Also, P_V and $F^1(P_{\omega,\alpha})$ define the same subspace modulo p . We shall identify $F^1(P_{\omega,\alpha})$ (respectively, $P_{\omega,\alpha}/F^1(P_{\omega,\alpha})$) with the direct sum of n copies (respectively, 1 copy) of A (via the natural identification, based on the definition of the Dieudonné modules of \mathcal{G}_Λ and $\mathbf{Q}_p/\mathbf{Z}_p$). Thus, by projection

$$P_F \hookrightarrow P_{\omega,\alpha} \rightarrow P_{\omega,\alpha}/F^1(P_{\omega,\alpha}) = A$$

we obtain a natural isomorphism of P_F with A , and, dually, a natural isomorphism of P_V with A^n (i.e., the direct sum of n copies of A). Finally, since

$$P_{\omega,\alpha} = P_F \oplus P_V$$

we may regard $F^1(P_{\omega,\alpha}) \subseteq P_{\omega,\alpha}$ as the graph of an A -linear morphism $A^n = P_V \rightarrow A = P_F$, which, by means of the various canonical trivializations, gives us an element

$$L_{\omega,z} \in p \cdot A^n$$

Lemma 2.13. *Suppose that $z \in U$, and let $\alpha \in (S_{\text{PD}})_z(A)$. Then we have*

$$L_{\omega,z} = \log_\Lambda(q_{\omega,z}(\alpha))$$

where “ \log_Λ ” is the morphism discussed in (5) of the discussion following Proposition 2.4.

Proof. When $n = 1$, $\lambda_1 = 1$, then this reduces to Theorem 1.7 of Chapter III of [Mzk1]. The proof here is entirely analogous: The point is to observe that the theory of [BK] allows one to interpret the morphism \log_Λ in the language of crystalline representations and their corresponding Dieudonné modules. The rest is an exercise in diagram-chasing. \circ

Now we return to the *compatibility of differentials* question. We no longer assume that $z \in U$. Then we have the following result:

Lemma 2.14. *We have $q_{\omega,z}^*(\omega_{\Lambda,\sigma}) = \pi_i^* \omega_\sigma$ (where $\sigma(1) = i$).*

Proof. When $n = 1$, $\lambda_1 = 1$, then this reduces to Theorem 1.8 of Chapter III of [Mzk1]. In *loc. cit.*, Theorem 1.8 (of *loc. cit.*) was deduced formally from Theorem 1.7 (of *loc. cit.*); one deduces Lemma 2.14 here from Lemma 2.13 above in an entirely analogous fashion. The idea (at least in the non-logarithmic case) is to interpret the coordinate morphism $q_{\omega,z}$ as being (essentially) the section of the projective bundle associated to $P_{\omega,\alpha} = P_F \oplus P_V$ defined by the Hodge filtration. Moreover, Lemma 2.13 tells us that, if we think of this projective bundle in terms of the coordinates defined by the canonical trivializations of P_F and P_V , then this section is the section defined by the function $\log_\Lambda(q_{\omega,\alpha})$. Taking derivatives thus shows that the derivative of this section (which is just the Kodaira-Spencer morphism of Hodge filtration of $P_{\omega,\alpha}$) is equal to the derivative of $\log_\Lambda(q_{\omega,\alpha})$. Tracing through the definitions shows that this is precisely what we wanted to prove. The logarithmic case is reduced to the non-logarithmic case by using the fact that the open subset of S where the log structure is trivial is dense in S , so any equality of logarithmic differentials on S is true if and only if it is true over a dense open subset (cf. [Mzk1], Chapter III, Theorem 1.8, as well as the discussion preceding Definition 1.11 for more details). \circ

Thus, in summary, we have the following result:

Theorem 2.15. *Suppose that we are given a binary-ordinary system of Frobenius liftings $\{\Phi_i^{\log}\}_{i=1,\dots,n}$ on S^{\log} . Let*

$$z \in S(k) \hookrightarrow S_{\text{PD}}(k)$$

where k is an algebraically closed field. Then for each $\omega \in \Omega_{\Phi_{\text{id}}}^{\text{et}}|_z$, there is a canonically associated coordinate morphism

$$q_{\omega,z} : (S_{\mathrm{PD}}^{\log})_z \rightarrow \mathcal{G}_{\Lambda}^{\log}$$

such that $q_{\omega,z}^*(\omega_{\Lambda,\sigma}) = \pi_i^* \omega_{\sigma}$ (where $\sigma(1) = i$).

In particular, if D is nonempty, so all the $\lambda_i = 1$, then we have a trace morphism

$$\mathrm{tr}_{\Lambda} : \mathcal{G}_{\Lambda} = \widehat{\mathbf{G}}_{\mathfrak{m}} \otimes_{\mathbf{Z}_p} \mathcal{O}_{\varpi} \rightarrow \widehat{\mathbf{G}}_{\mathfrak{m}}$$

induced by the trace $\mathcal{O}_{\varpi} \rightarrow \mathbf{Z}_p$. Let us consider the composite $\gamma_{\omega,z}$

$$S_z^{\log} \longrightarrow (S_{\mathrm{PD}}^{\log})_z \xrightarrow{q_{\omega,z}} \mathcal{G}_{\Lambda}^{\log} \xrightarrow{\mathrm{tr}_{\Lambda}} \widehat{\mathbf{G}}_{\mathfrak{m}}^{\log} \longrightarrow \mathbf{G}_{\mathfrak{m}}^{\log}$$

where the first morphism is the diagonal embedding, and the last morphism is the natural completion morphism. Suppose now that ω has residue e_j at the component of D defined by t_j . Then the pull-back via $\gamma_{\omega,z}$ of the divisor at infinity of $\mathbf{G}_{\mathfrak{m}}^{\log}$ has order at (t_j) equal to

$$\varpi \cdot e_j$$

(Note: the factor of ϖ arises from the fact that the pull-back via $\gamma_{\omega,z}$ of the divisor at infinity of $\mathbf{G}_{\mathfrak{m}}^{\log}$ is the restriction to $S \subseteq S_{\mathrm{PD}}$ of a union of divisors on each of the $n = \varpi$ factors of S in S_{PD} , each of which has order e_j at D .) Indeed, this follows immediately from Lemma 2.14 (cf. the discussion preceding Definition 1.11 in [Mzk1], Chapter III).

Now let us discuss the relationship between these canonical coordinates and the original system of Frobenius liftings.

Proposition 2.16. *The following diagram*

$$\begin{array}{ccc} (S_{\mathrm{PD}}^{\log})_z & \xrightarrow{q_{\omega,z}} & \mathcal{G}_{\Lambda}^{\log} \\ \downarrow (\Phi_{S_{\mathrm{PD}}})_z & & \downarrow p^n \\ (S_{\mathrm{PD}}^{\log})_z & \xrightarrow{q_{\omega,z}} & \mathcal{G}_{\Lambda}^{\log} \end{array}$$

commutes.

Proof. This follows from the commutative diagram of short exact sequences in the paragraph following Definition 2.10. \circ

Note that over $(S_{\mathrm{PD}})_z$, $G_{\Omega_{\Phi}}$ descends to a p -divisible group $\Theta_{\Phi_{\mathrm{id}}}^{\mathrm{et}} \otimes_{\mathcal{O}_{\varpi}} G_{\Lambda}$ over A . Let

$$\Theta_{\Lambda,z}$$

denote the corresponding formal group. Let us refer to the pull-back to S_{PD} of the tangent bundle of S_{FM}^{\log} (i.e., the dual of the bundle of differentials that are continuous with respect to the adic topology defined by diagonal) as the *tangent bundle of S_{PD}^{\log}* . Then we have the following rewording of Theorem 2.15 and Proposition 2.16:

Theorem 2.17. *Suppose that $z \in U(k)$. Then we have a commutative diagram*

$$\begin{array}{ccc} (S_{\text{PD}})_z & \xrightarrow{\Gamma_z} & \Theta_{\Lambda,z} \\ \downarrow (\Phi_{S_{\text{PD}}})_z & & \downarrow p^n \\ (S_{\text{PD}})_z & \xrightarrow{\Gamma_z} & \Theta_{\Lambda,z} \end{array}$$

where the horizontal morphisms Γ_z are those defined by the $q_{\omega,z}$ (as ω ranges over all elements of $\Omega_{\Phi_{\text{id}}}^{\text{et}}$). Moreover, Γ_z is formally étale in the following sense: it induces an isomorphism of the pull-back (via Γ_z) of the tangent bundle of $\Theta_{\Lambda,z}$ with the tangent bundle of $(S_{\text{PD}}^{\log})_z$.

We shall refer to Γ_z as the multi-uniformization via \mathcal{G}_{Λ} of $\Phi_{S_{\text{PD}}}$ at z .

Finally, before continuing, let us make the following definition/observation:

Definition 2.18. Let $z \in S(k)$. Then we shall call a point $\alpha \in S_{\text{PD}}(A)$ *multi-canonical* if its projection via π_i to a point in $S(A)$ is σ -canonical (where $\sigma(1) = i$), as in Definition 1.3.

Thus, every $z \in S(k)$ has a *unique* multi-canonical lifting to some

$$\alpha_z \in S_{\text{PD}}(A)$$

Moreover, we have the following result (which is obvious from the definitions and Lemma 2.13):

Proposition 2.19. *Let $z \in U(k)$; suppose that $\alpha \in S_{\text{PD}}(A)$ lifts z . Then α is multi-canonical if and only if for every $\omega \in \Omega_{\Phi_{\text{id}}}^{\text{et}}|_z$, $q_{\omega,z}(\alpha) = 1 \in \mathcal{G}_{\Lambda}(A)$.*

§2.4. Canonical Affine Coordinates

In this subsection, k need not be algebraically closed. Let

$$\alpha \in S_{\text{PD}}(A)$$

be *multi-canonical* and such that the reduction of α modulo p lies in $U(k)$. Let

$$\mathcal{A}^\alpha$$

be the p -adic completion of the PD-envelope of $\mathcal{O}_{S_{\text{PD}}}$ at the subscheme $\text{Im}(\alpha)$. Let

$$\epsilon_\alpha : \mathcal{A}^\alpha \rightarrow A$$

be the augmentation that defines the point α . Let

$$\mathcal{I} = \text{Ker}(\epsilon_\alpha)$$

The A -algebra structure, together with the augmentation ϵ_α define a splitting

$$\mathcal{A}^\alpha = A \oplus \mathcal{I}$$

which we shall call the *augmentation splitting* of \mathcal{A}^α . Note that we have an A -linear Frobenius action

$$\Phi_{\mathcal{A}} : (\mathcal{A}^\alpha)^{F^\varpi} \rightarrow \mathcal{A}^\alpha$$

induced by the Frobenius lifting $\Phi_{S_{\text{PD}}}$. Moreover, $\Phi_{\mathcal{A}}$ preserves the augmentation splitting, as well as the ring structure of \mathcal{A}^α .

Let us consider the *slopes* of this Frobenius action $\Phi_{\mathcal{A}}$. Clearly, $\Phi_{\mathcal{A}}$ acts on $A \oplus 0 \subseteq \mathcal{A}^\alpha$ as Φ_A^ϖ . Next, we note that since $\Phi_{\mathcal{A}}$ is manufactured out of a morphism on $\prod S$ which is a composite of n Frobenius liftings on each factor, it follows that $\Phi_{\mathcal{A}}$ maps \mathcal{I} into $p^n \cdot \mathcal{I}$. Thus, we have

$$\Phi_{\mathcal{A}}(\mathcal{I}^{[j]}) \subseteq p^{nj} \cdot \mathcal{I}^{[j]}$$

where the superscript in brackets denotes the divided power. By the definition of a binary-ordinary system of Frobenius liftings,

$$\Omega_\alpha \stackrel{\text{def}}{=} \mathcal{I}/\mathcal{I}^{[2]}$$

has *constant slope* n (with respect to the Frobenius action defined by $\Phi_{\mathcal{A}}$). Thus, if we divide $\Phi_{\mathcal{A}}$ restricted to Ω_α by p^n , we obtain an isomorphism

$$(\Omega_\alpha)^{F^\infty} \rightarrow \Omega_\alpha$$

which is the direct sum of the $\pi_i^* \Omega_{\Phi_\sigma}|_\alpha$ (where $\sigma(1) = i$).

Next, let us consider the A -submodule

$$\Omega^{\text{can}} \subseteq \mathcal{I}$$

which is the closure (in the p -adic topology) of the *intersection* of the images of

$$\left(\frac{1}{p^n} \cdot \Phi_{\mathcal{A}}|_{\mathcal{I}}\right)^N$$

in \mathcal{I} (for all $N \geq 1$). Since $\mathcal{I}/\mathcal{I}^{[2]}$ has slope n , it is clear that the projection

$$\Omega^{\text{can}} \rightarrow \Omega_\alpha$$

is *surjective*. Now let us consider the intersection

$$\Omega^{\text{can}} \cap \mathcal{I}^{[2]} \subseteq \mathcal{I}$$

Let $\phi = (\frac{1}{p^n} \cdot \Phi_{\mathcal{A}}|_{\mathcal{I}})^N(\psi)$, where $\psi \in \mathcal{I}$. If ϕ is contained in $\mathcal{I}^{[2]}$ modulo p^N , then since $\mathcal{I}/\mathcal{I}^{[2]}$ has slope n , it follows that ψ is also contained in $\mathcal{I}^{[2]}$ modulo p^N . But this implies that $\phi = (\frac{1}{p^n} \cdot \Phi_{\mathcal{A}}|_{\mathcal{I}})^N(\psi)$ is *zero* modulo p^N . Letting $N \rightarrow \infty$ thus shows that the intersection $\Omega^{\text{can}} \cap \mathcal{I}^{[2]}$ is *zero*. Thus, we conclude that the projection $\Omega^{\text{can}} \rightarrow \Omega_\alpha$ is an *isomorphism*. Inverting this isomorphism, we thus get a canonical morphism

$$\kappa_A : \Omega_\alpha \hookrightarrow \mathcal{A}^\alpha$$

Let

$$S_{\text{PD}}^\alpha \stackrel{\text{def}}{=} \text{Spf}(\mathcal{A}^\alpha)$$

(where the “Spf” is with respect to the p -adic topology). Let Θ_α be the dual A -module to Ω_α . Let

$$\Theta_\alpha^{\text{aff}}$$

be the p -adic completion of the PD-envelope at the origin of the $((nd)$ -dimensional) *affine space modeled on* Θ_α . Thus, $\Theta_\alpha^{\text{aff}}$ is Spf of the p -adic

completion of the PD-envelope of the symmetric algebra (over A) of Ω_α at its augmentation ideal. We may then reinterpret the canonical morphism κ_A as defining an *isomorphism*

$$\Delta^{\text{can}} : \Theta_\alpha^{\text{aff}} \cong S_{\text{PD}}^\alpha$$

We thus see that we have proven the following result:

Theorem 2.20. *For every choice of a multi-canonical*

$$\alpha \in S_{\text{PD}}(A)$$

at whose image the log structure of S_{PD}^{\log} is trivial, we obtain a local multi-uniformization (canonically associated to the given Frobenius system)

$$\Delta^{\text{can}} : \Theta_\alpha^{\text{aff}} \cong S_{\text{PD}}^\alpha$$

of S_{PD} at α by the affine space modeled on Θ_α .

Definition 2.21. We shall call the elements of the image of κ_A *canonical affine parameters associated to the Frobenius system at α* .

Finally, let us discuss the *relationship between these affine parameters and the group-theoretic coordinates of Theorem 2.15*. Let k be algebraically closed. Let $z \stackrel{\text{def}}{=} \alpha \pmod{p} \in S(k)$. Then for $\omega \in \Omega_{\Phi_{\text{id}}}^{\text{et}}$, we have

$$q_{\omega, \alpha} \stackrel{\text{def}}{=} q_{\omega, z} \in \mathcal{G}_\Lambda((S_{\text{PD}})_z)$$

(cf. Theorem 2.15). By taking the *logarithm* \log_Λ (cf. (5) of the discussion following Proposition 2.4) of $q_{\omega, \alpha}$, we obtain an element of $(\mathcal{A}^\alpha)^n$. If we add up the n components of this element, we thus get an element of \mathcal{A}^α which we denote by

$$\log_\Lambda^\Sigma(q_{\omega, \alpha}) \in \mathcal{A}^\alpha$$

On the other hand, the n -tuple given by $\{\omega_\sigma\}_\sigma$ defines an element of Ω_α . Thus, we also have an element $\kappa_A(\{\omega_\sigma\}_\sigma) \in \mathcal{A}^\alpha$.

Theorem 2.22. *Let k be algebraically closed. Let $z \stackrel{\text{def}}{=} \alpha \pmod{p} \in U(k) \subseteq S(k)$. Let $\omega \in \Omega_{\Phi_{\text{id}}}^{\text{et}}|_z$. Then we have*

$$\log_{\Lambda}^{\Sigma}(q_{\omega, \alpha}) = \kappa_A(\{\omega_{\sigma}\}_{\sigma})$$

in \mathcal{A}^{α} .

Proof. The proof is entirely similar to that of Theorem 1.15 of Chapter III of [Mzk1]. Indeed, one simply notes that both sides have the same derivative (cf. Lemma 2.14) and both vanish at α . \circ

§2.5. Lubin-Tate Geometries

In this subsection and the next, we pause to look at special cases of the theory developed so far in the hope of conveying its *geometric* meaning (as opposed to simply its technical meaning, in the form of such and such an isomorphism, etc.). For a more graphic representation of how the author envisions the Lubin-Tate and anabelian geometries, we refer the author to the Pictorial Appendix (as well as §1.6 of the Introduction to the present work).

In the present subsection, we consider the case $n = 1$, $\lambda_i = \lambda \geq 1$. For simplicity, we assume that the log structures are trivial. As the reader might expect, the assumption that $n = 1$ substantially simplifies the above theory. Thus, our initial data consists of a *single lifting*

$$\Phi : S \rightarrow S$$

of the λ^{th} power of Frobenius on $S_{\mathbf{F}_p}$ whose derivative divided by p

$$\frac{1}{p}d\Phi : \Phi^*\Omega_{S/A} \rightarrow \Omega_{S/A}$$

is an *isomorphism*. In the following, we shall call such a Frobenius lifting a *Lubin-Tate Frobenius lifting* (of order λ). Already, this looks much closer to the classical ordinary theory of [Mzk1], Chapter III, §1 (cf. Theorem 0.3 of the Introduction).

Although at first, if one takes a look back at the classical ordinary theory of [Mzk1], Chapter III, §1 (cf. Theorem 0.3 of the Introduction), it may seem that the “uniformization” obtained from a classical ordinary Frobenius lifting is given by a *special coordinate*, in fact, closer inspection reveals that really, the uniformization arising from a classical ordinary Frobenius lifting is best understood to be a *local isomorphism* between S and (a product of copies of) $\widehat{\mathbf{G}}_m$. Yet another interesting way to understand a classical ordinary uniformization is to regard it as the datum of a sort of *ring homomorphism*

$$M_d(\mathbf{Z}_p) \rightarrow \text{End}(S)$$

(to be understood in the spirit of *complex multiplication*), where d is the dimension of S over A , and we refrain from specifying precisely what we mean here by “End.” Thus, it is natural to ask:

If we consider spaces with complex multiplication by $M_d(\mathbf{Z}_p)$, why not consider spaces with complex multiplication by $M_d(\mathcal{O}_\lambda)$?

Such spaces would then naturally be modelled on the *Lubin-Tate formal group* \mathcal{G}_λ , and the uniformization would be given by a local isomorphism of S with (a product of copies of) \mathcal{G}_λ . This is precisely what we mean by a *Lubin-Tate uniformization*, or a *Lubin-Tate geometry*. Thus, in this language, Theorem 2.17 states precisely that a *Lubin-Tate Frobenius lifting gives rise to a Lubin-Tate geometry*.

Another reason that it is natural to expect a theory of Lubin-Tate geometry generalizing the classical ordinary theory of [Mzk1] is the following: The original motivating example for the theory of [Mzk1] was $\overline{\mathcal{M}}_{1,0}$, and its canonical Frobenius-invariant indigenous bundle. On the other hand, this phenomenon on $\overline{\mathcal{M}}_{1,0}$ is hardly unique: that is to say, this phenomenon is a special case of the *general theory of Shimura curves*. Moreover, in the theory of Shimura curves, there appear certain Shimura curves with Hecke-type correspondences that essentially compactify Lubin-Tate Frobenius liftings (just as $\overline{\mathcal{M}}_{1,0}$ admits a Hecke correspondence that compactifies its canonical ordinary Frobenius lifting) – cf. the theory of [Ih1-4], [Cara]. Thus, *given the existence of such Shimura curves, it is natural to expect a theory of Lubin-Tate uniformizations of $\overline{\mathcal{M}}_g$ as discussed in the present book*.

Finally, relative to the analogy (discussed in the Introduction of [Mzk1]) between classical ordinary Frobenius liftings and Kähler metrics, Lubin-Tate Frobenius liftings may, in a broad sense, be regarded as forming part of this analogy, except that there is nothing quite like a (nonclassical) Lubin-Tate Frobenius lifting over \mathbf{C} , since \mathbf{R} (unlike \mathbf{Q}_p) has no nontrivial unramified extensions. Or, put another way, *although there are many non-isomorphic formal groups (of the same dimension) over \mathbf{Z}_p , there is essentially only one (per dimension) metrized abelian complex Lie group germ over \mathbf{C}* .

§2.6. Anabelian Geometries

In this subsection, we discuss what happens when $n > 1$, but all the λ_i 's are equal (to some λ). This situation is a more radical departure from that of the classical ordinary case of [Mzk1], Chapter III, §1 (cf. Theorem 0.3 of the Introduction) than the Lubin-Tate case. That is to say, although the Lubin-Tate case still lies within the realm of p -adic objects analogous to a Kähler metric, in order to seek an analogue over \mathbf{C} of the case $n > 1$, one must leave the realm of classical Kähler geometry and consider various proposals for noncommutative geometry given by various authors (e.g., [Manin]). Although there are, in fact, lots of geometric situations over \mathbf{C} that go by the name “noncommutative geometry,” (e.g., quotients by foliations, “Spec”'s of noncommutative rings, etc.) these situations are not all equivalent to each other, and so one must be careful when speaking of “noncommutative geometry.” In the following, we attempt to explain what we mean when we say that a binary-ordinary Frobenius system of the type under consideration gives rise to a “noncommutative geometry.”

We begin by reviewing our data. For simplicity, we shall restrict ourselves to the non-logarithmic case. Then we assume that we are given n liftings (for $i = 1, \dots, n$)

$$\Phi_i : S \rightarrow S$$

of the λ^{th} power of Frobenius on $S_{\mathbf{F}_p}$. Moreover, we require that each Φ_i is (by itself) Lubin-Tate of order λ . We shall call such a system *anabelian of length n and order λ* . Already, one can see a nonabelian element here: in general,

The Φ_i will not commute with each other (see Lemma 2.24 below).

Thus, although each Φ_i alone defines a Lubin-Tate geometry, the *anabelian geometry* defined by the *system* will, in general, have little or nothing to do with those Lubin-Tate geometries.

To make things precise, we begin with a technical Lemma.

Lemma 2.23. *Then endomorphism ring $\text{End}(\mathcal{G}_\lambda)$ of the formal group \mathcal{G}_λ is equal to \mathcal{O}_λ . Similarly, for any integer $m \geq 1$, $\text{End}(\mathcal{G}_\lambda^m) = M_m(\mathcal{O}_\lambda)$.*

Proof. As noted previously, one has a natural morphism $\mathcal{O}_\lambda \hookrightarrow \text{End}(\mathcal{G}_\lambda)$. To see that this morphism is surjective, it suffices to consider endomorphisms of the Dieudonné module of \mathcal{G}_λ (cf. Example 2.7 of §2.1). The last statement follows formally from the first. \circ

Now we make precise first of all the observation that *it is very unusual for the Φ_i 's to commute with one another*. In fact, we have the following:

Lemma 2.24. *Suppose that Φ (respectively, Φ') is a Lubin-Tate Frobenius lifting on S of order λ (respectively, λ'). Suppose also that Φ commutes with Φ' . Then $\Phi = \Phi'$ (so, in particular, $\lambda = \lambda'$).*

Proof. Assume (without loss of generality) that k is algebraically closed. Let S_z^{PD} be the p -adic completion of the PD-envelope of S at some $z \in S(k)$. Write

$$S_z^{\text{PD}} = \text{Spf}(\mathcal{A}_z)$$

Observe that the canonical lifting $\alpha_z \in S(A)$ of z with respect to Φ is given by the p -adic limit

$$\alpha_z = \lim_{N \rightarrow \infty} \Phi^N(\beta)$$

(where $\beta \in S(A)$ is *any* lifting of z). By commutativity, it is thus a formal consequence of this formula for α_z that the *canonical liftings* of z with respect to Φ and Φ' *coincide*. Moreover, by considering constant slope subspaces of \mathcal{A}_z (as in Theorem 2.20), one sees immediately (by commutativity) that the *affine uniformizations* corresponding to Φ and Φ' *coincide*. It thus follows that Φ respects the group structure on S_z given by the uniformization of S_z by a product of G_{λ} 's arising from Φ' (as in Theorem 2.17). Indeed, to check this over S_z , it suffices to check it over S_z^{PD} since $\mathcal{O}_{S_z} \rightarrow \mathcal{O}_{S_z^{\text{PD}}}$ is *injective*; but over S_z^{PD} , this is clear, from our observation concerning affine uniformizations.

Thus, by Lemma 2.23, it follows that

$$\Phi|_{S_z} = [\xi]'$$

(where $\xi \in M_d(\mathcal{O}_{\lambda'})$, and $[-]'$ represents multiplication on S_z via the Lubin-Tate geometry arising from Φ'). Since $d\Phi$ is zero modulo p , ξ must be divisible by p ; since $d\Phi$ is only divisible *once* by p , it follows that $\xi = p \cdot u$ (where $u \in \text{GL}_d(\mathcal{O}_{\lambda'})$), hence that

$$\Phi|_{S_z} = [\xi]' = (\Phi'|_{S_z}) \circ [u]' = [u]' \circ (\Phi'|_{S_z})$$

Thus, since Φ and Φ' differ by composition with the *invertible* morphism $[u]'$, we conclude that Φ and Φ' have the same degree, so $\lambda = \lambda'$. Note that this implies that Φ and Φ' *coincide modulo p* .

Now, we use the fact that Φ and Φ' are equal modulo p . This implies that *the automorphism $[u]'$ of the formal group \mathcal{G}_{λ}^d is the identity modulo*

p . On the other hand, the reduction of \mathcal{G}_λ^d modulo p already determines the Dieudonné crystal of \mathcal{G}_λ^d (without its Hodge filtration); thus, we obtain that $[u]'$ induces the identity on the Dieudonné crystal of \mathcal{G}_λ^d . But this clearly means that $u = 1$, so $\Phi|_{S_z} = \Phi'|_{S_z}$. Since this equality holds for any $z \in S(k)$, the proof is complete. \square

Next, we would like to make precise the claim that in general,

$G_{\Phi_{PD}}$ does not descend to S_{FM} .

(Here, and in the following, by “descend” we shall mean “descend as an extension of $\mathbf{Q}_p/\mathbf{Z}_p$ by $G_{\Omega_\Phi|_{S_{FM}}}$.”) Let us concentrate, for simplicity, on the case $n = 2$. Then I claim that if $G_{\Phi_{PD}}$ did descend to some $G_{\Phi_{FM}}$ on S_{FM} , it would follow that there exists a Lubin-Tate Frobenius lifting

$$\Xi : S_{FM} \rightarrow S_{FM}$$

of order λ that induces a morphism

$$(\Phi_{(S_{FM})_{\mathbf{F}_p}}^\lambda)^* \mathbf{D}(G_{\Phi_{FM}}) = \Xi^* \mathbf{D}(G_{\Phi_{FM}}) \rightarrow \mathbf{D}(G_{\Phi_{FM}})$$

on the Dieudonné crystal $\mathbf{D}(G_{\Phi_{FM}})$ of $G_{\Phi_{FM}}$ that preserves the Hodge filtration. Indeed, one constructs Ξ modulo successive powers of p by deformation theory. The space of liftings of $\Xi_{\mathbf{Z}/p^{n+1}\mathbf{Z}}$ to $\mathbf{Z}/p^{n+2}\mathbf{Z}$ is a torsor over

$$(\Phi_{(S_{FM})_{\mathbf{F}_p}}^\lambda)^*(\Theta_{(S_{FM})_{\mathbf{F}_p}})$$

(where $\Theta_{(S_{FM})_{\mathbf{F}_p}}$ is the tangent bundle of $(S_{FM})_{\mathbf{F}_p}$ – cf. the discussion preceding Theorem 2.17) while the obstruction to preserving the Hodge filtration modulo p^{n+2} is a section of

$$\{F^1(\mathbf{D}(G_{\Phi_{FM}}))_{\mathbf{F}_p}\}^\vee = \Theta_{(S_{FM})_{\mathbf{F}_p}}$$

Moreover, one checks easily (by using the fact that the Kodaira-Spencer morphism for the Hodge filtration of $\mathbf{D}(G_{\Phi_{FM}})$ is the identity) that the morphism

$$(\Phi_{(S_{FM})_{\mathbf{F}_p}}^\lambda)^*(\Theta_{(S_{FM})_{\mathbf{F}_p}}) \rightarrow \Theta_{(S_{FM})_{\mathbf{F}_p}}$$

from deformations of Ξ to resulting obstructions to preserving the Hodge filtration is given by the Frobenius action on $F^1(\mathbf{D}(G_{\Omega_\Phi|_{S_{FM}}}))_{\mathbf{F}_p} = \Theta_{(S_{FM})_{\mathbf{F}_p}}$ (induced by the λ^{th} power of the natural Frobenius action on

the Dieudonné module $\mathbf{D}(G_{\Omega_{\Phi}})$). Thus, it follows that this morphism (from deformations to obstructions) is an *isomorphism*. In particular, we conclude that there always exists a unique deformation for which the corresponding obstruction is zero, i.e., for which the Hodge filtration is preserved. Finally, it is easy to compute that the derivative $\frac{1}{p} \cdot d\Xi$ is inverse to the above “deformations-to-obstructions morphism.” Thus, in particular, it follows that $\frac{1}{p} \cdot d\Xi$ is *invertible*. We refer to (the discussion preceding) [Mzk1], Chapter IV, Theorem 1.6, for a more detailed discussion of this sort of deformation argument.

Thus, to summarize, *under the assumption that $G_{\Phi_{\text{FM}}}$ descends to S_{FM}* , it follows that there exists a unique Lubin-Tate Frobenius lifting $\Xi : S_{\text{FM}} \rightarrow S_{\text{FM}}$ (of order λ) that preserves the Hodge filtration. But this would imply that, at $z \in S(k)$, we obtain a commutative diagram as in Theorem 2.17, except with S_{PD} (respectively, $\Phi_{S_{\text{PD}}}; p^n$) replaced by S_{FM} (respectively, $\Xi; p$). In particular, we obtain (by Theorem 2.17) that

$$\Xi^2 = \Phi_{S_{\text{FM}}}$$

Now let us write $\Psi_1 = (\Phi_1, \Phi_2) : S_{\text{FM}} \rightarrow S_{\text{FM}}$ (respectively, $\Psi_2 = (\Phi_2, \Phi_1) : S_{\text{FM}} \rightarrow S_{\text{FM}}$) for the morphism induced on S_{FM} by the product of Φ_1 and Φ_2 . Thus, the above equation implies (since, by definition, $\Phi_{S_{\text{FM}}} = \Psi_1 \circ \Psi_2$) that

$$\Xi^2 = \Psi_1 \circ \Psi_2$$

Now we have the following:

Lemma 2.25. *Suppose that Ψ_1 and Ψ_2 are Lubin-Tate Frobenius liftings of order λ on S_{FM} . Then the following three conditions are equivalent:*

(i) *There exists a Lubin-Tate Frobenius lifting Ξ of order λ on S_{FM} such that $\Psi_1 \circ \Psi_2 \equiv \Xi^2$ modulo p^3 .*

(ii) *There exists an automorphism Δ of S_{FM} which is the identity modulo p and which satisfies $\Psi_1 \circ \Delta^{-1} \equiv \Delta \circ \Psi_2$ modulo p^2 .*

(iii) *The two morphisms $\Phi_{(S_{\text{FM}})_{\mathbf{F}_p}}^* \Omega_{(S_{\text{FM}})_{\mathbf{F}_p}} \rightarrow \Omega_{(S_{\text{FM}})_{\mathbf{F}_p}}$ obtained by reducing $\frac{1}{p}(d\Psi_1)$ and $\frac{1}{p}(d\Psi_2)$ modulo p are equal.*

Proof. (iii) \implies (ii): If (iii) is satisfied, then Ψ_1 and Ψ_2 differ modulo p^2 by a section of

$$(\Phi_{(S_{\text{FM}})_{\mathbb{F}_p}}^\lambda)^{-1} \Theta_{(S_{\text{FM}})_{\mathbb{F}_p}} \subseteq (\Phi_{(S_{\text{FM}})_{\mathbb{F}_p}}^\lambda)^* \Theta_{(S_{\text{FM}})_{\mathbb{F}_p}}$$

(i.e., a section of $(\Phi_{(S_{\text{FM}})_{\mathbb{F}_p}}^\lambda)^* \Theta_{(S_{\text{FM}})_{\mathbb{F}_p}}$ obtained by pulling back a section of $\Theta_{(S_{\text{FM}})_{\mathbb{F}_p}}$ by $\Phi_{(S_{\text{FM}})_{\mathbb{F}_p}}^\lambda$). But this section clearly defines (modulo p^2) an automorphism Δ of S_{FM} which is the identity modulo p and which satisfies:

$$\Delta \circ \Psi_2 \equiv \Psi_1 \equiv \Psi_1 \circ \Delta^{-1} \quad (\text{modulo } p^2)$$

where the second congruence follows from the *general fact* that, *modulo* p^2 , any composite of the form

$$(\text{Frobenius lifting on } S_{\text{FM}}) \circ (\text{any morphism to } S_{\text{FM}})$$

depends only on the reduction of “any morphism” modulo p^2 . (Indeed, this follows from the fact that the morphism induced by a Frobenius lifting between tangent bundles modulo p is zero.)

(ii) \implies (i): If we let

$$\Xi \stackrel{\text{def}}{=} \Psi_1 \circ \Delta^{-1}$$

then the difference between Ξ^2 and

$$(\Psi_1 \circ \Delta^{-1}) \circ (\Delta \circ \Psi_2) = \Psi_1 \circ \Psi_2$$

modulo p^3 is given by a section

$$\alpha \in (\Phi_{(S_{\text{FM}})_{\mathbb{F}_p}}^{2\lambda})^* \Theta_{(S_{\text{FM}})_{\mathbb{F}_p}}$$

Moreover, α may be computed as the image in $(\Phi_{(S_{\text{FM}})_{\mathbb{F}_p}}^{2\lambda})^* \Theta_{(S_{\text{FM}})_{\mathbb{F}_p}}$ under the (pull-back by $\Phi_{(S_{\text{FM}})_{\mathbb{F}_p}}^\lambda$ of the) morphism on tangent bundles modulo p induced by $\Psi_1 \circ \Delta^{-1}$ of the section of

$$(\Phi_{(S_{\text{FM}})_{\mathbb{F}_p}}^\lambda)^* \Theta_{(S_{\text{FM}})_{\mathbb{F}_p}}$$

given by the difference between $\Xi = \Psi_1 \circ \Delta^{-1}$ and $\Delta \circ \Psi_2$ modulo p^3 . On the other hand, since this morphism on tangent bundles is zero, it follows that $\alpha = 0$, so $\Xi^2 = \Psi_1 \circ \Psi_2$ modulo p^3 , as desired.

(i) \implies (iii): First note that the fact that

$$\Psi_1 \circ \Psi_2 \equiv \Xi^2 \equiv \Xi \circ \Psi_2 \quad (\text{modulo } p^2)$$

(where the second “ \equiv ” follows from the “general fact” discussed in the first paragraph of this proof) *implies* (since Ψ_2 is faithfully flat) that $\Psi_1 \equiv \Xi$ modulo p^2 . Thus, $(\frac{1}{p} \cdot d\Psi_1) \equiv (\frac{1}{p} \cdot d\Xi)$ (modulo p). But, by (i), we have

$$(\frac{1}{p} \cdot d\Psi_2) \circ (\frac{1}{p} \cdot d\Psi_1) \equiv (\frac{1}{p} \cdot d\Xi) \circ (\frac{1}{p} \cdot d\Xi)$$

(modulo p). Since $\frac{1}{p} \cdot d\Psi_1$ is assumed to be invertible, this clearly implies (iii). \circ

In other words, by Lemma 2.25, we conclude that for an equation like $\Xi^2 = \Psi_1 \circ \Psi_2$ to hold even modulo p^3 , it would be necessary that Ω_{Φ_1} and Ω_{Φ_2} coincide modulo p . That is to say:

If Ω_{Φ_1} and Ω_{Φ_2} do not coincide modulo p , then $G_{\Phi_{\text{FM}}}$ does not descend to S_{FM} .

But it is easy to choose Φ_1 and Φ_2 so that Ω_{Φ_1} and Ω_{Φ_2} do not coincide modulo p . Thus, in general, $G_{\Phi_{\text{FM}}}$ does not descend to S_{FM} .

Let us return to the issue of understanding the geometry of an anabelian Frobenius system. In the Lubin-Tate case, the Lubin-Tate formal group \mathcal{G}_λ was a local model of the relevant geometry, i.e., the Lubin-Tate geometry. More generally, \mathcal{G}_Λ may be regarded as a sort of model for the geometry of a binary-ordinary Frobenius system of multi-order Λ . However, when $n > 1$, the *sense* in which \mathcal{G}_Λ is a *local* model changes: Namely, although by Theorems 2.17 and 2.20, it is certainly a local model when one “PD-localizes” to a sufficient extent, in general (as we just saw), it will *not* be a local model for $(S_{\text{FM}})_z$ (i.e., when one localizes by completing, but not by taking the PD-envelope). Put another way, one can see things most clearly if one PD-localizes “all the way,” as in Theorem 2.20, and in that case, the geometry defined by an anabelian system is just a usual affine geometry, but on a *product* of copies of S . (For that matter, if one PD-localizes a Lubin-Tate geometry, it looks pretty much the same as a classical ordinary geometry.) On the other hand, over $(S_{\text{FM}})_z$, or, *a fortiori*, over $S_z \subseteq (S_{\text{FM}})_z$, the various local geodesic coordinates or directions that exist PD-locally do not “split,” but are somehow tangled up together in some sort of *noncommuting fashion*. This is how the author envisions the geometry of an anabelian system.

One interesting observation is that in [Manin] and [Atiy], a proposal is given for “*new dimensions in arithmetic geometry*” that would generalize both the traditional geometric dimensions as well as the purely

arithmetic dimension of $\text{Spec}(\mathbf{Z})$. In particular, in [Atiy], it is proposed that these new dimensions on a variety M need not be entirely alien relative to the traditional geometric dimensions in that *they may be just the dimensions of some variety N containing M as a subvariety; nonetheless, N would be sharply peaked, or localized, at M .* It seems to the author that there is a certain *similarity* between what is proposed in [Manin] and [Atiy] and the geometry that actually occurs for an anabelian system: Namely, that sufficiently locally (i.e., PD-locally) this anabelian geometry is nothing particularly new, it is just a typical affine geometry on the larger object S_{PD} containing, and, sharply peaked at S . However, prior to this rather drastic (i.e., PD-) localization, the geodesics of this familiar affine geometry are twisted or tangled up in some complicated noncommutative way inside $(S_{\text{FM}})_z$ or S_z .

§2.7. Deformation of the System of Frobenius Liftings

The purpose of this subsection is to prove Lemma 2.11 of §2.2. We begin by considering a certain concrete example, which will be of importance in the proof of Lemma 2.11. In this subsection, k (a perfect field of characteristic p), n , and $\lambda_1, \dots, \lambda_n$ (natural numbers) will be as in §2.2.

Let

$$R \stackrel{\text{def}}{=} k(x_i^{\pm \frac{1}{p^\infty}})_{i=1, \dots, n}$$

where the x_i (for $i = 1, \dots, n$) are indeterminates, and the symbol “ $\pm \frac{1}{p^\infty}$ ” in the exponent of x_i means that we adjoin all p -power roots of x_i and x_i^{-1} to k . Thus, R is a perfect field. Let us denote by

$$X_i \stackrel{\text{def}}{=} [x_i] \in W(R)$$

the Teichmüller representative of x_i .

Next, let

$$\mathcal{O}_S \stackrel{\text{def}}{=} W(R)[[T]]$$

where T is an indeterminate. Let us write $t \in \mathcal{O}_S \otimes \mathbf{F}_p$ for the image of T in $\mathcal{O}_S \otimes \mathbf{F}_p$. Write $S \stackrel{\text{def}}{=} \text{Spf}(\mathcal{O}_S)$ (where the “Spf” is with respect to the p -adic topology on R). Let

$$\Phi_i : S \rightarrow S$$

be the morphism defined by

$$\Phi_i^{-1}(T) = X_i \cdot T + T^{p^{\lambda_i}}$$

Thus, if we set

$$\Xi_i \stackrel{\text{def}}{=} X_i + T^{p^{\lambda_i}-1} \in \mathcal{O}_S$$

(and write $\xi_i \in \mathcal{O}_S \otimes \mathbf{F}_p$ for the reduction of Ξ_i modulo p), then

$$\frac{1}{p} \cdot d\Phi_i(dT) = \Xi_i \cdot dT$$

In particular, since Ξ_i is an invertible element of \mathcal{O}_S , it follows that Φ_1, \dots, Φ_n define a *binary-ordinary system of Frobenius liftings* (cf. Definition 2.1 – the only difference between our discussion here and the situation of Definition 2.1 is that \mathcal{O}_S is complete in the “geometric direction,” but this difference is inessential relative to what we wish to do in the present subsection). If (for σ a shifting permutation) we define $\Phi_\sigma \stackrel{\text{def}}{=} \Phi_{\sigma(1)} \circ \dots \circ \Phi_{\sigma(n)}$ (cf. §1.1), then

$$\frac{1}{p^n} \cdot d\Phi_\sigma(dT) = \Xi_\sigma \cdot dT$$

for some invertible element $\Xi_\sigma \in \mathcal{O}_S$. If we write $\xi_\sigma \in \mathcal{O}_S \otimes \mathbf{F}_p$ for the reduction of Ξ_σ modulo p , then we have

$$\xi_\sigma = \xi_{\sigma(n)} \cdot \xi_{\sigma(n-1)}^{p^{\lambda_{\sigma(n)}}} \cdot \xi_{\sigma(n-2)}^{p^{\lambda_{\sigma(n-1)} + \lambda_{\sigma(n)}}} \cdot \dots \cdot \xi_{\sigma(1)}^{p^{\lambda_{\sigma(2)} + \lambda_{\sigma(3)} + \dots + \lambda_{\sigma(n)}}}$$

Let

$$\mathcal{O}_{S'} \stackrel{\text{def}}{=} \mathcal{O}_S [\Xi_i^{\frac{-1}{p^{\varpi}-1}}]_{i=1, \dots, n}$$

(where $\varpi = \lambda_1 + \dots + \lambda_n$). Thus, if we write $S' \stackrel{\text{def}}{=} \text{Spf}(\mathcal{O}'_S)$ (where the “Spf” is with respect to the p -adic topology), then S' is a *connected finite étale cover* of S . Moreover, we have elements

$$\zeta_\sigma \stackrel{\text{def}}{=} \xi_\sigma^{\frac{-1}{p^{\varpi}-1}} \in \mathcal{O}_{S'} \otimes \mathbf{F}_p$$

Let Σ be the set of all the *shifting permutations* (cf. §1.1). Then we have the following result:

Lemma 2.26. *The n elements*

$$\{\zeta_\sigma\}_{\sigma \in \Sigma} \subseteq \mathcal{O}_{S'} \otimes \mathbf{F}_p$$

are linearly independent over $\overline{\mathbf{F}}_p$.

Proof. First, observe that it suffices to show that the ζ_σ are linearly independent over $\overline{\mathbf{F}}_p$ in the R -subalgebra $C \subseteq \mathcal{O}_{S'} \otimes \mathbf{F}_p$ generated by t and the ζ_σ 's. Let $B \stackrel{\text{def}}{=} R[t] \subseteq C$. Now suppose that some nonzero $\overline{\mathbf{F}}_p$ -linear combination of these elements is equal to 0. Let ζ_σ be a term (among those terms with nonzero coefficients) for which $\lambda_{\sigma(1)}$ is minimal. Note that the ξ_i (for $i = 1, \dots, n$) are relatively prime in B . Indeed, the zeroes of ξ_i are the $(p^{\lambda_i} - 1)^{\text{th}}$ roots of x_i ; thus, distinct ξ_i have distinct zero loci. Now the choice of σ implies that for any $\sigma' \neq \sigma$ (such that the coefficient of $\zeta_{\sigma'}$ is nonzero), and every $i \geq 2$,

$$\begin{aligned} \lambda_{\sigma(2)} + \dots + \lambda_{\sigma(n)} = \varpi - \lambda_{\sigma(1)} &\geq \varpi - \lambda_{\sigma'(1)} = \lambda_{\sigma'(2)} + \dots + \lambda_{\sigma'(n)} \\ &> \lambda_{\sigma'(i+1)} + \dots + \lambda_{\sigma'(n)} \end{aligned}$$

(where, in the last inequality, we use the fact that all the λ_j 's are > 0). In particular, it follows that the zero of ξ_σ at the roots of $x_{\sigma(1)}$ (which will be of order $\lambda_{\sigma(2)} + \dots + \lambda_{\sigma(n)}$) will be of strictly greater order than the zero of any $\xi_{\sigma'}$ at those points (which will be of order $\lambda_{\sigma'(i+1)} + \dots + \lambda_{\sigma'(n)}$, for that i – necessarily ≥ 2 (since $\sigma'(1) \neq \sigma(1)$) – such that $\sigma'(i) = \sigma(1)$). Thus, the pole of ζ_σ at the roots of $x_{\sigma(1)}$ will be of strictly greater order than the pole of any $\zeta_{\sigma'}$ at those points. This contradicts the assertion that ζ_σ is a linear combination of such $\zeta_{\sigma'}$, as desired. \circ

Next, let us observe that the \mathbf{F}_{p^ϖ} -subspace of $\Omega_{S_{\mathbf{F}_p}}$ (where the “ Ω ” denotes continuous differentials with respect to the topology of \mathcal{O}_S defined by the maximal ideal of \mathcal{O}_S) given by $\Omega_{\Phi_\sigma}^{\text{et}} \otimes \mathbf{F}_p$ is generated by $\zeta_\sigma \cdot dT$. (Indeed, this follows from the definition of Ξ_σ .) That is to say, we have

$$\Omega_{\Phi_\sigma}^{\text{et}} \otimes \mathbf{F}_p = \mathbf{F}_{p^\varpi} \cdot (\zeta_\sigma \cdot dT) \subseteq \Omega_{S_{\mathbf{F}_p}}$$

In particular, interpreting Lemma 2.26 in terms of the morphism κ^{et} of §2.2 gives the following:

Lemma 2.27. *The morphism*

$$\kappa^{\text{et}} : \bigoplus_{\sigma} (\Omega_{\Phi_\sigma}^{\text{et}} \otimes \mathbf{F}_p) \rightarrow \Omega_{S_{\mathbf{F}_p}^{\log}}$$

discussed in §2.2 in injective.

Proof. Indeed, the injectivity of κ^{et} follows from the linear independence discussed in Lemma 2.26. \circ

Now we are ready to prove Lemma 2.11 of §2.2. First observe that Lemma 2.27 shows that (by taking a product of d copies of the Frobenius liftings constructed above) there exists a system of binary-ordinary Frobenius liftings (with the same n ; $\lambda_1, \dots, \lambda_n$ as the given system) for which κ^{et} is injective. Suppose that this system of Frobenius liftings is defined on some

$$W(\tilde{R})[[T_1, \dots, T_d]]$$

where \tilde{R} is a perfect field containing the original k , and T_1, \dots, T_d are indeterminates. Then by working over some localization \mathcal{R} (containing the points $\epsilon = 0$ and $\epsilon = 1$) of the perfection of the ring

$$\tilde{R}[\epsilon]$$

we can construct a system of Frobenius liftings on

$$W(\mathcal{R})[[T_1, \dots, T_d]]$$

whose power series are given by

$$\begin{aligned} &\epsilon \cdot (\text{the power series of the original system}) \\ &+ (1 - \epsilon) \cdot (\text{the power series of the system for which } \kappa^{\text{et}} \text{ is injective}) \end{aligned}$$

This new “ ϵ -family” of Frobenius liftings is such that at $\epsilon = 0$, its corresponding κ^{et} is *injective*. Moreover, it gives rise to a morphism “ β_m ” between finite étale covers of

$$\text{Spec}(W(\mathcal{R})[[T_1, \dots, T_d]] \otimes \mathbf{Q}_p)$$

(Of course, strictly speaking, we only defined “ β_m ” in the case when “ \mathcal{R} ” is a field, but it is easy to see that one can define an analogous morphism “ β_m ” in the present situation.) Finally, since $W(\mathcal{R})[[T_1, \dots, T_d]]$ is a *domain*, it follows that

$$\text{Spec}(W(\mathcal{R})[[T_1, \dots, T_d]] \otimes \mathbf{Q}_p)$$

is *irreducible*, as desired. This completes the proof of Lemma 2.11, and hence also of Theorem 2.12 of §2.2.

§3. Application to Curves and their Moduli

In this §, we briefly note how the general theory of the preceding two §'s can be applied in the case of hyperbolic curves and their moduli stacks.

§3.1. Frobenius Liftings on the Moduli Stack

Let g, r be nonnegative integers such that $2g - 2 + r \geq 1$. Let Π be a *binary VF-pattern* (i.e., such that $\Pi(\mathbf{Z}) \subseteq \{0, \chi\}$, where $\chi = \frac{1}{2}(2g - 2 + r)$) of period ϖ . Let

$$S^{\log} = ((\overline{\mathcal{N}}_{g,r}^{\Pi,s})^{\log})_{\mathbf{Z}_p}^{\text{ord}}$$

(cf. Chapter VII, Definition 1.3). Thus, S^{\log} is a p -adic formal log stack whose underlying stack is formally étale over $(\overline{\mathcal{M}}_{g,r})_{\mathbf{Z}_p}$ and formally smooth over \mathbf{Z}_p . Let

$$X^{\log} \rightarrow S^{\log}$$

be the *tautological curve* over S^{\log} . For each *active* $j \in \mathbf{Z}$ (i.e., j such that $\Pi(j) \neq 0$), we have (according to Chapter VII, Theorem 1.8) a *canonical Frobenius lifting*

$$\Phi_j^{\log} : S^{\log} \rightarrow S^{\log}$$

Let $0 = j_0, j_1, \dots, j_n = \varpi$ be the active integers between 0 and ϖ . We define the ordered set Λ as follows:

$$\Lambda = \{\lambda_1, \dots, \lambda_n\} \stackrel{\text{def}}{=} \{j_1 - j_0, j_2 - j_1, \dots, j_n - j_{n-1}\}$$

Then by Chapter VII, Theorem 3.2, it follows that the $\{\Phi_j^{\log}\}$ form a *binary-ordinary Frobenius system of multi-order Λ* . We state this as a Theorem:

Theorem 3.1. *Let*

$$S^{\log} = ((\overline{\mathcal{N}}_{g,r}^{\Pi,s})^{\log})_{\mathbf{Z}_p}^{\text{ord}}$$

where $2g - 2 + r > 0$ and Π is binary. Then the associated canonical system of Frobenius liftings

$$\Phi_j^{\log} : S^{\log} \rightarrow S^{\log}$$

(for active j) of Chapter VII, Theorem 1.8, is binary-ordinary in the sense of Definition 2.1.

In particular, one obtains on S^{\log} (and hence, étale locally on $(\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{Z}_p}$) group-theoretic multi-uniformizations as in Theorems 2.15 and 2.17, multi-canonical points as in Definition 2.18, and affine multi-uniformizations as in Theorem 2.20. For instance, if Π is of pure tone ϖ , then one obtains a canonical Lubin-Tate geometry on S^{\log} . If Π is pre-home, then one obtains a canonical anabelian geometry of length ϖ (and order 1) on S^{\log} .

Remark. Let us consider the case where $g = r = 1$, $p = 5$, and Π is pre-home of period 2. Let $S \stackrel{\text{def}}{=} (\overline{\mathcal{N}}_{1,1}^{\Pi, \text{ord}})_{\mathbf{Z}_p}$. Then by Theorem 3.1, we obtain a pair

$$\Phi_1, \Phi_2 : S \rightarrow S$$

of Frobenius liftings on S . We claim that over some nonempty open subset of $S_{\mathbf{F}_p}$, the morphisms $(\Omega_{\Phi_1})_{\mathbf{F}_p}$ and $(\Omega_{\Phi_2})_{\mathbf{F}_p}$ are distinct. In other words, relative to the discussion following Lemma 2.25 (in §2.6), the truth of this claim means that the geometry defined by the system consisting of this Φ_1^{\log} and Φ_2^{\log} is “truly anabelian.”

Let $T \stackrel{\text{def}}{=} (\overline{\mathcal{N}}_{1,1}^{\text{ord}})_{\mathbf{Z}_p}$, and let

$$\Psi : T \rightarrow T$$

be the canonical classical ordinary Frobenius lifting of [Mzk1], Chapter III, Theorem 2.8. Note that it follows from the definition of a “VF-stack” (cf. Chapter III, Definition 1.10) that

$$S_{\mathbf{F}_p} = T_{\mathbf{F}_p} \times_{(\overline{\mathcal{M}}_{1,1})_{\mathbf{F}_p}} T_{\mathbf{F}_p}$$

Thus, to see that $(\Omega_{\Phi_1})_{\mathbf{F}_p}$ and $(\Omega_{\Phi_2})_{\mathbf{F}_p}$ are distinct over some nonempty open subset of $S_{\mathbf{F}_p}$, it suffices to see that the morphism $(\Omega_{\Psi})_{\mathbf{F}_p}$ does not (generically) descend from $(\overline{\mathcal{N}}_{1,1}^{\text{ord}})_{\mathbf{F}_p}$ to $(\overline{\mathcal{M}}_{1,1}^{\text{ord}})_{\mathbf{F}_p}$. (Indeed, once one knows that there exist two distinct points in some fiber of $(\overline{\mathcal{N}}_{1,1}^{\text{ord}})_{\mathbf{F}_p} = T_{\mathbf{F}_p} \rightarrow (\overline{\mathcal{M}}_{1,1}^{\text{ord}})_{\mathbf{F}_p}$ at which $(\Omega_{\Psi})_{\mathbf{F}_p}$ takes distinct values, it follows that these two points define a point of $S_{\mathbf{F}_p}$ at which $(\Omega_{\Phi_1})_{\mathbf{F}_p}$ and $(\Omega_{\Phi_2})_{\mathbf{F}_p}$ “are” the two distinct values of $(\Omega_{\Psi})_{\mathbf{F}_p}$.)

But suppose it did descend. Since we know that the natural morphism $(\overline{\mathcal{N}}_{1,1}^{\text{ord}})_{\mathbb{F}_p} \rightarrow (\overline{\mathcal{M}}_{1,1})_{\mathbb{F}_p}$ is *surjective* in this case (cf. the second Remark following Theorem 1.4 of Chapter IV), this would imply – by regarding $(\Omega_\Psi)_{\mathbb{F}_p}$ as an everywhere nonvanishing section of

$$\text{Hom}(\Phi^*_{(\overline{\mathcal{M}}_{1,1})_{\mathbb{F}_p}} \omega_{(\overline{\mathcal{M}}_{1,1})_{\mathbb{F}_p}}^{\log}, \omega_{(\overline{\mathcal{M}}_{1,1})_{\mathbb{F}_p}}^{\log}) = \omega_{(\overline{\mathcal{M}}_{1,1})_{\mathbb{F}_p}}^{\otimes(1-p)}$$

(where the “Hom” is “sheaf Hom”) – that we have a modular form on $(\overline{\mathcal{M}}_{1,1})_{\mathbb{F}_p}$ of positive weight that vanishes nowhere, which is clearly absurd (since $\omega_{(\overline{\mathcal{M}}_{1,1})_{\mathbb{F}_p}}^{\log}$ is *ample* on $(\overline{\mathcal{M}}_{1,1})_{\mathbb{F}_p}$). This contradiction completes the proof of the claim.

§3.2. Frobenius Liftings on the Universal Curve

We maintain the notation of §3.1. Recall the *canonical indigenous bundles*

$$\mathcal{P}_j$$

(for active j) on X^{\log} (cf. Chapter VII, Theorem 1.8). Suppose that (i, j) is Π -adjacent. Note that there exists a morphism

$$h_j : (\Phi_X^*)^{j-i} \tau_{X_{\mathbb{F}_p}^{\log}/S_{\mathbb{F}_p}^{\log}} \rightarrow \text{Ad}(\mathcal{P}_j)_{\mathbb{F}_p} \rightarrow \tau_{X_{\mathbb{F}_p}^{\log}/S_{\mathbb{F}_p}^{\log}}$$

given by composing the $(j - i - 1)^{\text{st}}$ Frobenius pull-back of the p -curvature of $\mathbf{F}^*(\mathcal{P}_i)_{\mathbb{F}_p}$ with the projection to the quotient F^{-1}/F^0 of the Hodge filtration of $\text{Ad}(\mathcal{P}_j)_{\mathbb{F}_p}$. Let

$$(X^{\log})^{\text{ord}} \subseteq X^{\log}$$

be the p -adic formal open substack over which all the h_j are *isomorphisms* of line bundles. Note that $(X^{\log})^{\text{ord}}$ is dense in every fiber of $X^{\log} \rightarrow S^{\log}$.

Now by deformation theory, the fact that the Kodaira-Spencer morphism of an indigenous bundle is the identity, and the fact that the h_j are isomorphisms over $(X^{\log})^{\text{ord}}$, it follows that *there exist unique Frobenius liftings*

$$\Psi_j^{\log} : (X^{\log})^{\text{ord}} \rightarrow (X^{\log})^{\text{ord}}$$

that cover $\Phi_j^{\log} : S^{\log} \rightarrow S^{\log}$ and which have the property that the induced Frobenius action

$$(\Psi_j^{\log})^* \mathrm{Ad}(\mathcal{P}_i) \rightarrow \mathrm{Ad}(\mathcal{P}_j) \otimes \mathbf{Q}_p$$

preserves the Hodge filtration. We refer to [Mzk1], Chapter IV, Theorem 1.6; [Mzk1], Chapter V, Theorem 2.6, for more detailed discussions of this construction in the classical ordinary case. We also remark in passing that precisely the same sort of deformation argument (i.e., to construct a unique Frobenius lifting that preserves the Hodge filtration of some sort of uniformizing \mathcal{MF}^∇ -object) was used to construct the morphism “ Ξ ” in the discussion preceding Lemma 2.25 of §2.6.

We shall call these liftings Ψ_j^{\log} *canonical*. Moreover, these Frobenius liftings form a binary-ordinary system of multi-order Λ and period ϖ (cf. [Mzk1], Chapter IV, Theorem 1.6; [Mzk1], Chapter V, Theorem 2.6, for a more detailed discussion of the classical ordinary case). Thus, we have the following result:

Theorem 3.2. *Notation as in Theorem 3.1. Then there is a unique system of Frobenius liftings*


$$\Psi_j^{\log} : (X^{\log})^{\mathrm{ord}} \rightarrow (X^{\log})^{\mathrm{ord}}$$

that cover $\Phi_j^{\log} : S^{\log} \rightarrow S^{\log}$ and which have the property that the induced Frobenius action $(\Psi_j^{\log})^* \mathrm{Ad}(\mathcal{P}_i) \rightarrow \mathrm{Ad}(\mathcal{P}_j) \otimes \mathbf{Q}_p$ preserves the Hodge filtration. Moreover, this system is binary-ordinary of multi-order Λ and period ϖ .



In particular, one obtains on $(X^{\log})^{\mathrm{ord}}$ (and hence, étale locally on the universal curve $\mathcal{C}_{\mathbf{Z}_p}^{\log}$ over $(\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{Z}_p}$) group-theoretic multi-uniformizations as in Theorems 2.15 and 2.17, multi-canonical points as in Definition 2.18, and affine multi-uniformizations as in Theorem 2.20. For instance, if Π is of pure tone ϖ , then one obtains a canonical Lubin-Tate geometry on $(X^{\log})^{\mathrm{ord}}$. If Π is pre-home, then one obtains a canonical anabelian geometry of length ϖ (and order 1) on $(X^{\log})^{\mathrm{ord}}$.

Pictorial Appendix


We begin by giving pictures of *Lubin-Tate geometries* on a curve, which we envision as abelian affine geometries of “different colors.” First, we have the tone $\varpi = 1$ case, i.e., the *classical ordinary case*:



Then we have various *higher tone* (i.e., $\varpi > 1$) cases:



Finally, we give an illustration of the *anabelian case* (for $n = 2$), where the geometry on the curve is derived from an abelian affine geometry on a product of the curve with itself:



Chapter IX: The Geometrization of Spiked Frobenius Liftings

§0. Introduction

In this Chapter, we continue our discussion of the *geometrization of Frobenius liftings*, begun in Chapter VIII. This time, we deal with *very ordinary spiked Frobenius liftings* (of period 2) on a log scheme S^{\log} . Moreover, we use this geometrization to show how one can equip the *mantle* of S^{\log} (i.e., the PD-envelope of the diagonal of $S^{\log} \times S^{\log}$) with the structure of an \mathcal{MF}^{∇} -object in such a way that it gives rise to a *Galois representation* (Theorem 2.11). In fact, we will develop this “mantle theory” for the binary-ordinary case (Definition 2.13) as well as the very ordinary spiked case (Definition 2.12).

The key point in the spiked case is to note that the zeroes in the p -curvature give rise to certain *strong modular variables* and a virtual fibration from the moduli stack down to a space of strong modular variables. These strong modular variables carry a natural *Lubin-Tate geometry of order 2* (Theorem 1.5). Then, over the space of strong modular variables, the fibers of the moduli stack carry a natural *anabelian geometry of length 2 and order 1* (Theorem 2.3). (We refer to Chapter VIII, §2.6, for a detailed discussion of such geometries.) Thus, it is as though each renormalized Frobenius pull-back that one must visit upon the canonical indigenous bundle “anabelianizes” the geometry by one step. Of course, these two phenomena, i.e., that

- (1) each *zero* in the p -curvature “*Lubin-Tate-izes*” the geometry by a step;
- (2) each *active integer* in a period of Π “*anabelianizes*” the geometry by a step;

can be seen in an isolated setting in the binary-ordinary case (Chapter VIII). Indeed, this is why we presented the binary-ordinary case *first*.

What makes the spiked case special, though, is that in the spiked case, these two phenomena occur in a *mixed* fashion, thus giving rise to an interesting and intricate p -adic geometry (étale locally) on the moduli stack of curves (Theorem 3.1) – as well as on the universal curve (Theorem 3.4) – that is different in many respects from the simpler classical ordinary or binary-ordinary geometries studied thus far.

§1. The Formal Uniformizing \mathcal{MF}^∇ -Object

§1.1. The Objects in Question

Let k be a *perfect field of odd characteristic p* that contains \mathbf{F}_{p^2} . Let $A = W(k)$, the ring of Witt vectors with coefficients in k ; let K be its quotient field. Let S be a formally smooth, geometrically connected p -adic formal scheme over A of constant relative dimension d . Let S^{\log} be a log formal scheme whose underlying formal scheme is S , and whose log structure is given by a relative divisor with normal crossings $D \subseteq S$ over A . Let $\Phi_A : A \rightarrow A$ be the Frobenius morphism on A . Let us denote the result of base changing by Φ_A by means of a superscripted “ F .” In this Chapter, we would like to consider a lifting

$$\Phi^{\log} : S^{\log} \rightarrow S^{\log}$$

of the *square* of $\Phi_{S_{\mathbf{F}_p}} : S_{\mathbf{F}_p} \rightarrow S_{\mathbf{F}_p}$ (the Frobenius morphism on $S_{\mathbf{F}_p}$). Let us write

$$\Omega_\Phi \stackrel{\text{def}}{=} \frac{1}{p} \cdot d\Phi^{\log} : \Phi^* \Omega_{S^{\log}} \rightarrow \Omega_{S^{\log}}$$

and

$$\Omega''_{S_{\mathbf{F}_p}^{\log}} \stackrel{\text{def}}{=} \text{Ker}\{(\Omega_\Phi)_{\mathbf{F}_p} : (\Phi_{S_{\mathbf{F}_p}}^2)^* \Omega_{S_{\mathbf{F}_p}^{\log}} \rightarrow \Omega_{S_{\mathbf{F}_p}^{\log}}\} \subseteq (\Phi_{S_{\mathbf{F}_p}}^2)^* \Omega_{S_{\mathbf{F}_p}^{\log}} = \Phi^* \Omega_{S_{\mathbf{F}_p}^{\log}}$$

Next, let us denote by

$$\Omega'_{S^{\log}} \subseteq (\Phi^* \Omega_{S^{\log}}) \otimes_{\mathbf{Z}_p} \left(\frac{1}{p} \cdot \mathbf{Z}_p\right)$$

the subsheaf of sections whose reductions modulo p lie in $\Omega''_{S_{\mathbf{F}_p}^{\log}} \otimes_{\mathbf{Z}_p} (\frac{1}{p} \cdot \mathbf{Z}_p)$. (Here, $\Omega_{S^{\log}}$ denotes the sheaf of p -adically continuous differentials of S^{\log} over A .) Then it follows that Ω_Φ induces a morphism

$$\Omega'_\Phi : \Omega'_{S^{\log}} \rightarrow \Omega_{S^{\log}}$$

Now let us make the following assumptions:

- (1) $\Omega''_{S^{\log}_{\mathbf{F}_p}}$ is a *vector bundle* of rank $d - c$ on $S_{\mathbf{F}_p}$. Here we call the integer c the *colevel* (of Φ).
- (2) $d\Phi^{\log}$ induces *multiplication by p^2 on residues* of differential 1-forms at components of D .
- (3) Ω'_Φ is an *isomorphism*.
- (4) The *image* of the composite $(\Omega_\Phi \circ \Phi^*(\Omega_\Phi))_{\mathbf{F}_p}$ is a *subbundle* of $\Omega_{S^{\log}_{\mathbf{F}_p}}$ of rank c .

These assumptions may be interpreted more intuitively as follows: *Condition (1)* states that the *dimension* of the portion of $\Omega_{S^{\log}}$ on which the Frobenius action Ω_Φ has a *given slope* is *locally constant* on S . *Condition (2)* states that the Frobenius action Ω_Φ is compatible with (i.e., “ \mathbf{G}_m -like” for) the log structure of S^{\log} . *Condition (3)* will – in the context of canonical modular Frobenius liftings (as in Theorem 3.1) – amount to the condition of “ Π -ordinariness” studied in Chapter VII (cf. Chapter VII, Definition 1.1). *Condition (4)* will – in the context of canonical modular Frobenius liftings (as in Theorem 3.1) – correspond to (part of) the condition of “*very Π -ordinariness*” studied in Chapter VII (cf. Chapter VII, Definition 3.6, (b.)). Note that (assuming that Condition (1) is in force) Condition (4) is equivalent to the condition that for each point $s \in S$,

$$\text{Im}(\Phi^*(\Omega_\Phi)_{\mathbf{F}_p})_s \cap (\Omega''_{S^{\log}})_s = 0$$

(where “Im” stands for “the image of the morphism”). Now we are ready to make the following:

Definition 1.1. We shall say that a Frobenius lifting $\Phi^{\log} : S^{\log} \rightarrow S^{\log}$ satisfying (1) through (4) above is a *very ordinary spiked Frobenius lifting* (of colevel c).

§1.2. The Strong Portion of the Uniformization

We maintain the notation of the §1.1, and assume, moreover, that we are dealing with a *very ordinary spiked Frobenius lifting* $\Phi^{\log} : S^{\log} \rightarrow S^{\log}$. Let us consider the morphism

$$\Omega_\Phi : \Phi^* \Omega_{S^{\log}} \rightarrow \Omega_{S^{\log}}$$

Note that by Condition (4) of Definition 1.1 (for $N \geq 1$) the image of Ω_{Φ}^N modulo p is a vector subbundle of $\Omega_{S_{\mathbb{F}_p}}^{\log}$ of rank c which is *independent* of N . Similarly, by induction, one checks that given an integer $i \geq 1$, there exists an integer N_i such that the image of Ω_{Φ}^N modulo p^i is independent of $N \geq N_i$. Thus, by taking the *p-adic closure of the intersection of the images of Ω_{Φ}^N* (for $N \geq 1$), we see that we obtain a subbundle

$$\Omega_{S^{\log}}^{\text{st}} \subseteq \Omega_{S^{\log}}$$

of rank c which is stabilized by Ω_{Φ} and is such that Ω_{Φ} induces an isomorphism

$$\Omega_{\Phi}^{\text{st}} : \Phi^* \Omega_{S^{\log}}^{\text{st}} \cong \Omega_{S^{\log}}^{\text{st}}$$

Let

$$\Omega_{S^{\log}}^{\text{wk}} \stackrel{\text{def}}{=} (\Omega_{S^{\log}}) / (\Omega_{S^{\log}}^{\text{st}})$$

Thus, $\Omega_{S^{\log}}^{\text{wk}}$ is a *vector bundle of rank $d - c$* (where $d = \dim_A(S)$) on S . Moreover, the endomorphism induced by Ω_{Φ} on the quotient $\Omega_{S^{\log}}^{\text{wk}}$ is clearly divisible by p ; thus, we denote the *isomorphism* (cf. Condition (3) of Definition 1.1) induced by $\frac{1}{p} \cdot \Omega_{\Phi}$ by

$$\Omega_{\Phi}^{\text{wk}} : \Phi^* \Omega_{S^{\log}}^{\text{wk}} \cong \Omega_{S^{\log}}^{\text{wk}}$$

Moreover, we have the following:

Lemma 1.2. *The residues of differentials in $\Omega_{S^{\log}}^{\text{st}} \subseteq \Omega_{S^{\log}}$ at the irreducible components of the divisor $D \subseteq S$ are zero. Thus, in the future, we will often just write Ω_S^{st} for $\Omega_{S^{\log}}^{\text{st}}$.*

Proof. By Condition (2) of Definition 1.1, $d\Phi^{\log}$ acts with slope 2 on the residues at components of D . On the other hand, by definition, $d\Phi^{\log}$ acts with slope 1 on $\Omega_{S^{\log}}^{\text{st}}$. This completes the proof. \circ

As usual, we can form the $\Omega_{\Phi}^{\text{st}}$ -invariants (respectively, $\Omega_{\Phi}^{\text{wk}}$ -invariants) of Ω_S^{st} (respectively, $\Omega_{S^{\log}}^{\text{wk}}$); this gives rise to *local systems*

$$\Omega_S^{\text{st,et}} \quad \text{and} \quad \Omega_{S^{\log}}^{\text{wk,et}}$$

of free $W(\mathbb{F}_{p^2})$ -modules of rank c and $d - c$, respectively, on S .

Definition 1.3. We shall call the local system $\Omega_S^{\text{st}, \text{et}}$ (respectively, $\Theta_S^{\text{st}, \text{et}}$, i.e., the dual of $\Omega_S^{\text{st}, \text{et}}$) the *strong differential (tangential) local system associated to the Frobenius lifting*. Similarly, we shall call the local system $\Omega_{S^{\log}}^{\text{wk}, \text{et}}$ (respectively, $\Theta_{S^{\log}}^{\text{wk}, \text{et}}$, i.e., the dual of $\Omega_{S^{\log}}^{\text{wk}, \text{et}}$) the *weak differential (tangential) local system associated to the Frobenius lifting*.

Moreover, as usual, we can *integrate the differentials of $\Omega_S^{\text{st}, \text{et}}$* by defining an \mathcal{MF}^∇ -object (cf. Chapter VIII, §2.1, 2.2), as follows. Let

$$P_0^{\text{st}} \stackrel{\text{def}}{=} \Omega_S^{\text{st}} \oplus \mathcal{O}_S$$

Let $\nabla_{P_0^{\text{st}}}$ be the *connection* on P_0^{st} obtained (as usual) by first taking the *direct sum* of the connection on Ω_S^{st} for which sections of $\Omega_S^{\text{st}, \text{et}}$ are horizontal and the trivial connection on \mathcal{O}_S , and *adding* to this direct sum connection the $\text{End}(P_0^{\text{st}})$ -valued *differential*

$$P_0^{\text{st}} \rightarrow \Omega_S^{\text{st}} = (0 \oplus \mathcal{O}_S) \otimes_{\mathcal{O}_S} \Omega_S^{\text{st}} \hookrightarrow P_0^{\text{st}} \otimes_{\mathcal{O}_S} \Omega_{S^{\log}}$$

where the first morphism is the natural projection morphism, and the second morphism is the natural inclusion. Let

$$F^0(P_0^{\text{st}}) \stackrel{\text{def}}{=} P_0^{\text{st}}; \quad F^1(P_0^{\text{st}}) \stackrel{\text{def}}{=} \Omega_S^{\text{st}} \oplus 0; \quad F^2(P_0^{\text{st}}) \stackrel{\text{def}}{=} 0$$

Thus, the *Kodaira-Spencer morphism* for $F^1(-)$ is precisely the dual of the *inclusion $\Omega_S^{\text{st}} \hookrightarrow \Omega_{S^{\log}}$* . Let

$$(P_1^{\text{st}}, \nabla_{P_1^{\text{st}}}) \stackrel{\text{def}}{=} \mathbf{F}^*(P_0^{\text{st}}, \nabla_{P_0^{\text{st}}})$$

where the “ \mathbf{F}^* ” is defined (as usual) by considering sections whose reductions modulo p lie in $F^1(-)$. We endow the vector bundle P_1^{st} with the following *Hodge filtration*:

$$F^0(P_1^{\text{st}}) = P_1^{\text{st}}; \quad F^1(P_1^{\text{st}}) = 0$$

Finally, we define

$$\Phi_{P^{\text{st}}, 10} : \Phi_{S_{\mathbf{F}_p}}^*(P_1^{\text{st}}, \nabla_{P_1^{\text{st}}}) \cong (P_0^{\text{st}}, \nabla_{P_0^{\text{st}}})$$

(where $\Phi_{S_{\mathbf{F}_p}}^*$ is the “pull-back on crystals via the Frobenius morphism on $S_{\mathbf{F}_p}$ ”) as follows: Note that the vector bundle on the *left-hand side* is just

$$\Phi^*(\Omega_S^{\text{st}} \oplus \mathcal{O}_S)$$

while on the *right*, it is just

$$\Omega_S^{\text{st}} \oplus \mathcal{O}_S$$

Thus, by taking the direct sum of Ω_Φ^{st} with the identity on \mathcal{O}_S , we obtain an isomorphism of the sort desired.

Now note that we have horizontal injections $\iota_0 : \mathcal{O}_S \hookrightarrow P_0^{\text{st}}$ and $\iota_1 : \mathcal{O}_S \hookrightarrow P_1^{\text{st}}$ (where the second is the Frobenius pull-back of the first). Now we let $(P^{\text{st}}, \nabla_{P^{\text{st}}})$ be the *inductive limit* of the following diagram:

$$\begin{array}{ccc} \mathcal{O}_S \oplus \mathcal{O}_S & \xrightarrow{\iota_0 \oplus \iota_1} & (P_0^{\text{st}}, \nabla_{P_0^{\text{st}}}) \oplus (P_1^{\text{st}}, \nabla_{P_1^{\text{st}}}) \\ \downarrow & & \\ \mathcal{O}_S & & \end{array}$$

where the vertical morphism is the identity on both factors of \mathcal{O}_S . We endow P^{st} with the *Hodge filtration* induced from that of the P_i^{st} , and finally, we note that $\Phi_{P^{\text{st}}, 10}$ and $(P_1^{\text{st}}, \nabla_{P_1^{\text{st}}}) \stackrel{\text{def}}{=} \mathbf{F}^*(P_0^{\text{st}}, \nabla_{P_0^{\text{st}}})$ define a *Frobenius action* $\Phi_{P^{\text{st}}}$ on $(P^{\text{st}}, \nabla_{P^{\text{st}}})$. Thus, in summary, we have defined an \mathcal{MF}^∇ -object

$$(P^{\text{st}}, \nabla_{P^{\text{st}}}, F^\cdot(P^{\text{st}}), \Phi_{P^{\text{st}}})$$

Moreover, this \mathcal{MF}^∇ -object defines a *p-divisible group* G_Φ^{st} on S which fits into an *exact sequence of p-divisible groups*

$$0 \rightarrow (\Theta_S^{\text{st}, \text{et}}) \otimes_{W(\mathbf{F}_{p^2})} (G_2|_S) \rightarrow G_\Phi^{\text{st}} \rightarrow \mathbf{Q}_p/\mathbf{Z}_p \rightarrow 0$$

where (cf. Chapter VIII, §2.1 – especially Example 2.7) G_2 is the *Lubin-Tate p-divisible group associated to* $W(\mathbf{F}_{p^2})$.

Let

$$\Pi_{S^{\log}} \stackrel{\text{def}}{=} \pi_1(S_K^{\log}, s)$$

(where we choose some basepoint $s \in S(K)$ at which the log structure is trivial). Then by restricting the *p-divisible group* G_Φ^{st} to S_K^{\log} , and applying $\text{Hom}(-, \mathbf{Q}_p/\mathbf{Z}_p)$, we obtain a $\Pi_{S^{\log}}$ -*module* $G_{\text{Gal}}^{\text{st}}$, which fits into an *exact sequence of $\Pi_{S^{\log}}$ -modules*

$$0 \rightarrow \mathbf{Z}_p \rightarrow G_{\text{Gal}}^{\text{st}} \rightarrow (\Omega_S^{\text{st}, \text{et}}) \otimes_{W(\mathbf{F}_{p^2})} (T_p(G_2)) \rightarrow 0$$

where $T_p(G_2)$ is the p -adic Tate module associated to the p -divisible group G_2 .

Definition 1.4. We shall call G_{Φ}^{st} (respectively, $G_{\text{Gal}}^{\text{st}}$) the *strong uniformizing p -divisible group* (respectively, *Galois representation*) associated to the Frobenius lifting.

Moreover, by applying Kummer theory for the formal group \mathcal{G}_2 (cf. Chapter VIII, Proposition 2.6) to the exact sequence of p -divisible groups just constructed, we obtain the following result (cf. Chapter VIII, Theorem 2.17):

Theorem 1.5. Suppose that $z \in S(k)$, where k is algebraically closed. Let

$$\Theta_z^{\text{st}}$$

be the formal group over A corresponding to $(\Theta_S^{\text{st,et}})|_z \otimes_{W(\mathbf{F}_{p,2})} (G_2)$. Then we have a commutative diagram

$$\begin{array}{ccc} S_z & \xrightarrow{\Gamma_z^{\text{st}}} & \Theta_z^{\text{st}} \\ \downarrow \Phi_z^{\log} & & \downarrow p \\ S_z & \xrightarrow{\Gamma_z^{\text{st}}} & \Theta_z^{\text{st}} \end{array}$$

where Γ_z^{st} is formally smooth, and the morphism induced by Γ_z^{st} on differentials is precisely the natural inclusion

$$\Omega_S^{\text{st,et}}|_z \hookrightarrow \Omega_{S_z}$$

We shall refer to Γ_z^{st} as the strong uniformization (via \mathcal{G}_2) at z (defined by the Frobenius lifting).

Finally, before proceeding, we make the following construction:
Let

$$S^{\text{st}}[n] \rightarrow S$$

be the finite flat morphism of degree p^{2nc} parametrizing *splittings* of the above exact sequence of p -divisible groups (i.e., preceding Definition 1.4) over $(\frac{1}{p^n}\mathbf{Z}_p)/\mathbf{Z}_p \subseteq \mathbf{Q}_p/\mathbf{Z}_p$ (cf. the discussion preceding Chapter VIII, Definition 2.9). Thus, we obtain a projective system

$$\dots \rightarrow S^{\mathrm{st}}[n+1] \rightarrow S^{\mathrm{st}}[n] \rightarrow \dots \rightarrow S^{\mathrm{st}}[0] = S$$

We shall denote the *first morphism* $S^{\mathrm{st}}[1] \rightarrow S^{\mathrm{st}}[0]$ of this system by

$$T \rightarrow S$$

Note that it follows immediately from Theorem 1.5 that all of the $S^{\mathrm{st}}[n]$ are *formally smooth p -adic formal schemes over A* and that the pull-back to $S^{\mathrm{st}}[n]$ of the divisor $D \subseteq S$ is a *divisor with normal crossings on $S^{\mathrm{st}}[n]$* , hence defines a *log structure on $S^{\mathrm{st}}[n]$* .

§1.3. The Strong Portion of the Mantle

Let us denote the *PD-envelope of the diagonal* $S^{\mathrm{log}} \hookrightarrow S^{\mathrm{log}} \times_A S^{\mathrm{log}}$ by:

$$\mathcal{A}$$

We regard \mathcal{A} as an \mathcal{O}_S -algebra via the structure from the *right*. There is a natural logarithmic connection $\nabla_{\mathcal{A}}$ on \mathcal{A} for which sections of \mathcal{O}_S from the *left* are *horizontal*. Let

$$\mathcal{I} \subseteq \mathcal{A}$$

be the PD-ideal defining the diagonal. Thus, $\mathcal{A}/\mathcal{I} = \mathcal{O}_S$. We shall regard \mathcal{A} as filtered via the “*Hodge filtration*”

$$F^i(\mathcal{A}) = \mathcal{I}^{[i]}$$

(where the superscript in brackets is a divided power). In particular, $F^1/F^2(\mathcal{A}) = \Omega_{S^{\mathrm{log}}}$. We shall refer to

$$(\mathcal{A}, \nabla_{\mathcal{A}}, F^{\bullet}(\mathcal{A}))$$

as the *mantle of S^{log}* . Finally, we equip the mantle with a *Frobenius action*

$$\Phi_{\mathcal{A}} : (\Phi_{S_{\mathbb{F}_p}}^*)^2(\mathcal{A}, \nabla_{\mathcal{A}}) = \Phi^*(\mathcal{A}, \nabla_{\mathcal{A}}) \rightarrow (\mathcal{A}, \nabla_{\mathcal{A}})$$

given by applying Φ^{log} to the factor of S^{log} on the *left*. Of course, $\Phi_{\mathcal{A}}$ is not an isomorphism (in general).

Note that we have a horizontal inclusion $\iota : \mathcal{O}_S \hookrightarrow \mathcal{A}$ from the *right*. Let

$$\mathcal{D} \stackrel{\text{def}}{=} \mathcal{A}/\mathcal{O}_S$$

Thus, \mathcal{D} inherits a *connection* $\nabla_{\mathcal{D}}$ from \mathcal{A} . Moreover, $\Phi_{\mathcal{A}}$ descends to a *Frobenius action* $\Phi_{\mathcal{D}}$ on \mathcal{D} . Since $\Phi_{\mathcal{A}}$ annihilates \mathcal{I} modulo p , it follows that $\Phi_{\mathcal{D}}$ is *divisible* by p .

Let us consider the p -adic closure

$$\mathcal{D}^1 \subseteq \mathcal{D}$$

of the intersection of the images of the $(\frac{1}{p} \cdot \Phi_{\mathcal{D}})^N$ (for $N \geq 1$). Since $\Phi_{\mathcal{D}}$ is compatible with $\nabla_{\mathcal{D}}$, it follows that $\nabla_{\mathcal{D}}$ stabilizes \mathcal{D}^1 .

Lemma 1.6. *We have a natural horizontal isomorphism of vector bundles on S :*

$$\mathcal{D}^1 \cong \Omega_S^{\text{st}}$$

(where the connection on Ω_S^{st} is the one for which sections of $\Omega_S^{\text{st,et}}$ are horizontal) that is compatible with the Frobenius actions given by $\frac{1}{p} \cdot \Phi_{\mathcal{D}}$ and $\Omega_{\Phi}^{\text{st}}$.

Proof. To compute $\Phi_{\mathcal{D}}$, it suffices to consider $\Phi_{\mathcal{A}}$, where we regard $(\Phi_{S_{\mathbb{F}_p}}^*)^2 \mathcal{A}$ as $\Phi^* \mathcal{A}$. Then $\Phi_{\mathcal{A}}$ becomes the morphism $\mathcal{A} \rightarrow \mathcal{A}$ given by applying Φ^{\log} on both factors. This morphism *stabilizes* \mathcal{I} . Let us denote the induced morphism by

$$\Phi_{\mathcal{I}} : \Phi^* \mathcal{I} \rightarrow \mathcal{I}$$

Moreover, the composite

$$\delta : \mathcal{I} \hookrightarrow \mathcal{A} \rightarrow \mathcal{D}$$

is an *isomorphism* compatible with the respective Frobenius actions (i.e., $\Phi_{\mathcal{I}}$ and $\Phi_{\mathcal{D}}$). But $\mathcal{I}/\mathcal{I}^2 = \Omega_{S^{\log}}$, and the action of $\frac{1}{p} \cdot \Phi_{\mathcal{I}}$ on \mathcal{I}^2 is divisible by p . Thus, it follows that $\delta^{-1}(\mathcal{D}^1) \cap \mathcal{I}^2 = 0$ (cf. the discussion of Chapter VIII, §2.4), so we have an inclusion

$$\mathcal{D}^1 \hookrightarrow \Omega_{S^{\log}}$$

which is compatible with the Frobenius actions (induced by $\frac{1}{p} \cdot \Phi_{\mathcal{D}}$ on the left, and Ω_{Φ} on the right).

Moreover, the image of this inclusion is the slope 0 (with respect to Ω_{Φ}) portion of $\Omega_{S^{\log}}$. Indeed, this follows from the fact that if we lift a

Frobenius (i.e., $\frac{1}{p} \cdot \Phi_{\mathcal{I}}$) invariant section of $\mathcal{I}/\mathcal{I}^2 = \Omega_{S^{\log}}$ to a section s of \mathcal{I} , then as $N \rightarrow \infty$, the section $(\frac{1}{p} \cdot \Phi_{\mathcal{I}})^N(s)$ of \mathcal{I} converges p -adically (since the action of $\frac{1}{p} \cdot \Phi_{\mathcal{I}}$ on \mathcal{I}^2 is divisible by p) to some s_{∞} which is $\frac{1}{p} \cdot \Phi_{\mathcal{I}}$ invariant and has the same image as s in $\mathcal{I}/\mathcal{I}^2$. This shows that we get an isomorphism $\mathcal{D}^1 \cong \Omega_S^{\text{st}}$ which is compatible with the respective Frobenius actions.

The fact that this isomorphism is compatible with the connections follows from the fact that since the Frobenius action on \mathcal{D}^1 is compatible with its connection, any Frobenius invariant section d of \mathcal{D}^1 (i.e., any section of $\Omega_S^{\text{st,et}} \subseteq \Omega_S^{\text{st}} \cong \mathcal{D}^1$) is necessarily horizontal. (Indeed,

$$\nabla_{\mathcal{D}}(d) = \nabla_{\mathcal{D}}(\Phi_{\mathcal{D}}^N(d)) = \Phi_{\mathcal{D}}^N(\nabla_{\mathcal{D}}(d)) \in p^N \cdot (\mathcal{D}^1 \otimes_{\mathcal{O}_S} \Omega_{S^{\log}})$$

for all $N \geq 1$, so $\nabla_{\mathcal{D}}(d) = 0$.) \circ

Next, let us define $\text{Geo}^{\text{st}}(\mathcal{A})$ by the fiber product:

$$\text{Geo}^{\text{st}}(\mathcal{A}) \stackrel{\text{def}}{=} \mathcal{A} \times_{\mathcal{D}} \mathcal{D}^1 \subseteq \mathcal{A}$$

Note that $\text{Geo}^{\text{st}}(\mathcal{A})$ is preserved by $\nabla_{\mathcal{A}}$ and $\Phi_{\mathcal{A}}$. Moreover, $\text{Geo}^{\text{st}}(\mathcal{A})$ is a vector bundle of rank $1+c$, consisting of functions (i.e., sections of \mathcal{A}) of the form:

$$(\text{constant term}) + (\text{strong linear terms})$$

By definition, $\text{Geo}^{\text{st}}(\mathcal{A})$ fits into an exact sequence

$$0 \rightarrow \mathcal{O}_S \rightarrow \text{Geo}^{\text{st}}(\mathcal{A}) \rightarrow \Omega_S^{\text{st}} \rightarrow 0$$

stabilized by $\nabla_{\mathcal{A}}$ and $\Phi_{\mathcal{A}}$.

In general, whenever we have a vector bundle \mathcal{E} on S which fits into an exact sequence

$$0 \rightarrow \mathcal{O}_S \rightarrow \mathcal{E} \rightarrow \mathcal{V} \rightarrow 0$$

(where \mathcal{V} is also a vector bundle), we define the *affinization* of \mathcal{E} (relative to this exact sequence) as the quasi-coherent \mathcal{O}_S -algebra

$$\text{Aff}(\mathcal{E})$$

whose “Spf” is the \mathcal{V}^{\vee} -torsor that parametrizes splittings of the above exact sequence. Thus, $\text{Aff}(\mathcal{E})$ is locally (non-canonically) isomorphic

to (the p -adic completion of) a polynomial algebra over \mathcal{O}_S in $\text{rank}(\mathcal{V})$ variables. Moreover, the \mathcal{O}_S -submodule of $\text{Aff}(\mathcal{E})$ that corresponds to the polynomials of degree ≤ 1 is naturally isomorphic to \mathcal{E} itself.

Definition 1.7. We shall refer to $\text{Aff}(\mathcal{E})$ as the *affinization* of \mathcal{E} (with respect to the above exact sequence).

Let us return to our discussion of $\text{Geo}^{\text{st}}(\mathcal{A})$. Note that the \mathcal{O}_S -subalgebra of \mathcal{A} generated by $\text{Geo}^{\text{st}}(\mathcal{A})$ may be identified with $\text{Aff}(\text{Geo}^{\text{st}}(\mathcal{A}))$, i.e., we have a natural inclusion

$$\text{Aff}(\text{Geo}^{\text{st}}(\mathcal{A})) \hookrightarrow \mathcal{A}$$

Now we define

$$\text{Aff}(\text{Geo}^{\text{st}}(\mathcal{A})) \subseteq \mathcal{A}^{\text{st}} \subseteq \mathcal{A}$$

to be the p -adic completion of the PD-envelope of $\text{Aff}(\text{Geo}^{\text{st}}(\mathcal{A}))$ with respect to the augmentation $\text{Aff}(\text{Geo}^{\text{st}}(\mathcal{A})) \hookrightarrow \mathcal{A} \rightarrow \mathcal{A}/\mathcal{I} = \mathcal{O}_S$ defined by the diagonal embedding. Note that \mathcal{A}^{st} is *preserved* by $\nabla_{\mathcal{A}}$ and $\Phi_{\mathcal{A}}$. Thus, \mathcal{A}^{st} admits a *connection* $\nabla_{\mathcal{A}^{\text{st}}} \stackrel{\text{def}}{=} \nabla_{\mathcal{A}}|_{\mathcal{A}^{\text{st}}}$, as well as a *Hodge filtration*

$$F^{\cdot}(\mathcal{A}^{\text{st}}) \stackrel{\text{def}}{=} F^{\cdot}(\mathcal{A}) \bigcap \mathcal{A}^{\text{st}}$$

Definition 1.8. We shall refer to $\mathcal{A}^{\text{st}} \subseteq \mathcal{A}$ as the *strong portion* of the mantle \mathcal{A} .

§1.4. The Renormalized Frobenius Pull-back of the Mantle

In this subsection, we forget about very ordinary spiked Frobenius liftings, and discuss a certain *fundamental construction concerning the mantle* of S^{\log} .

The *mantle* $(\mathcal{A}, \nabla_{\mathcal{A}})$ with its natural connection forms a *crystal* on $\text{Crys}(S_{\mathbf{F}_p}^{\log})$. Thus, we can consider its “*naive Frobenius pull-back*”

$$(\mathcal{F}, \nabla_{\mathcal{F}}) \stackrel{\text{def}}{=} \Phi_{S_{\mathbf{F}_p}}^*(\mathcal{A}, \nabla_{\mathcal{A}})$$

Thus, \mathcal{F} is a p -adic quasi-coherent sheaf of \mathcal{O}_S -algebras. Note that there exists a morphism

$$\lambda : \Phi_{S_{\mathbf{F}_p}}^{-1} \mathcal{O}_S \rightarrow \mathcal{F}$$

(where $\Phi_{S_{\mathbf{F}_p}}^{-1}$ indicates the sheaf-theoretic pull-back on the Zariski site) arising from the inclusion $\mathcal{O}_S \hookrightarrow \mathcal{A}$ from the left.

Let

$$\beta : F \rightarrow S$$

be the *structure morphism* of $\mathrm{Spf}(\mathcal{F})$ (where we write “Spf” rather than “Spec” because we are working *p-adically*). Let

$$\gamma \stackrel{\mathrm{def}}{=} \mathrm{Spf}(\lambda) : F \rightarrow S$$

Let us endow F with the *log structure pulled back via β* . Thus, β extends to a log morphism $\beta^{\mathrm{log}} : F^{\mathrm{log}} \rightarrow S^{\mathrm{log}}$. Note that, with this log structure on F , γ also extends naturally to a morphism $\gamma^{\mathrm{log}} : F^{\mathrm{log}} \rightarrow S^{\mathrm{log}}$. Indeed, this follows from the fact that if t is a local equation for an irreducible component of D on S then the elements $t \otimes 1$ and $1 \otimes t$ (in \mathcal{A}) are *unit* multiples of one another in \mathcal{A} (cf. the definition of the “log PD-envelope” – [Kato], §5).

Note that we may regard F as defining the *functor* (which, by abuse of notation, we shall also denote by “ F ”) on \mathbf{Z}_p -flat S^{log} -log schemes U^{log} given by

$$U^{\mathrm{log}} \mapsto \{\text{morphisms } U \rightarrow F \text{ over } S\} = \{\text{morphisms } U^{\mathrm{log}} \rightarrow F^{\mathrm{log}} \text{ over } S^{\mathrm{log}}\}$$

Now let us consider the *subfunctor* $F' \subseteq F$ consisting of those morphisms $\zeta : U \rightarrow F$ such that the *reduction modulo p* of the *composite*

$$\gamma_U \stackrel{\mathrm{def}}{=} \gamma \circ \zeta : U \rightarrow S$$

is equal to the composite of the *structure morphism* $U_{\mathbf{F}_p} \rightarrow S_{\mathbf{F}_p}$ of U with the *Frobenius morphism* $\Phi_{S_{\mathbf{F}_p}} : S_{\mathbf{F}_p} \rightarrow S_{\mathbf{F}_p}$. Another way to phrase this condition is to say that the morphism

$$\zeta^{-1} : \Phi_{S_{\mathbf{F}_p}}^* \mathcal{A} \rightarrow \mathcal{O}_U$$

(arising from ζ) *vanishes* on the ideal $\Phi_{S_{\mathbf{F}_p}}^* \mathcal{I} \subseteq \Phi_{S_{\mathbf{F}_p}}^* \mathcal{A}$ defining the *diagonal*.

Observe that such a morphism $\gamma_U \stackrel{\mathrm{def}}{=} \gamma \circ \zeta : U \rightarrow S$ always extends naturally to a *log morphism* γ_U^{log} : Indeed, ζ always extends naturally to a log morphism $\zeta^{\mathrm{log}} : U^{\mathrm{log}} \rightarrow F^{\mathrm{log}}$ (by the definition of the log structure on F^{log}), so we just compose ζ^{log} with γ^{log} . We shall always regard F'

as an S -functor via the restriction to $F' \subseteq F$ of the structure morphism $\beta : F \rightarrow S$; denote the resulting morphism

$$\beta' \stackrel{\text{def}}{=} \beta|_{F'} : F' \rightarrow S$$

On the other hand, we also have:

$$\gamma' \stackrel{\text{def}}{=} \gamma|_{F'} : F' \rightarrow S$$

Lemma 1.9. *The S -functor F' is representable by a smooth (hence, in particular, \mathbf{Z}_p -flat) p -adic scheme over S whose reduction modulo p is equal to the torsor of Frobenius liftings (as in the discussion preceding Proposition 2.1 of [Mzk1], Chapter III).*

Proof. If X^{\log} is a (p -adic formal) *log scheme*, we shall denote its log structure by

$$\mu_X : M_X \rightarrow \mathcal{O}_X$$

and write

$$\Gamma_X \stackrel{\text{def}}{=} M_X / \mathcal{O}_X^\times$$

Note that the natural morphism

$$\Gamma_X \rightarrow \Gamma_{X_{\mathbf{F}_p}}$$

is an *isomorphism*. If δ is a section of Γ_X , we shall denote by

$$\mathcal{L}_\delta$$

the \mathcal{O}_X^\times -torsor which is the inverse image of $\delta \in \Gamma_X$ under the projection $M_X \rightarrow \Gamma_X$.

It suffices to show the representability of F' *étale locally* on S . Thus, we may assume that there exists a *lifting*

$$\tilde{\Phi}^{\log} : S^{\log} \rightarrow S^{\log}$$

of the Frobenius morphism $\Phi_{\mathbf{F}_p}^{\log} : S_{\mathbf{F}_p}^{\log} \rightarrow S_{\mathbf{F}_p}^{\log}$, and that $D \subseteq S$ is a union of smooth divisors. We may also assume that there is a collection of functions t_1, \dots, t_d on S such that

- (1) t_1, \dots, t_i define the *irreducible components* of $D \subseteq S$, while t_{i+1}, \dots, t_d are *units*;
- (2) $d \log(t_1), \dots, d \log(t_d)$ form a *basis* of $\Omega_{S^{\log}}$.

Since $\mu_S : M_S \hookrightarrow S$ is *injective* and its image contains the t_j , it follows that the t_j define unique sections

$$\log(t_j)$$

of M_S . Let

$$\log^\Gamma(t_j)$$

denote the image of $\log(t_j)$ in Γ_S .

Suppose that we are given $\zeta : U \rightarrow F$ and form the resulting $\gamma_U^{\log} : U^{\log} \rightarrow S^{\log}$. Let

$$\tilde{\Phi}_U^{\log} : U^{\log} \rightarrow S^{\log}$$

be the *composite* of the *structure morphism* $U^{\log} \rightarrow S^{\log}$ of U^{\log} with the *Frobenius lifting* $\tilde{\Phi}^{\log}$. Let

$$\mathcal{M}_j \stackrel{\text{def}}{=} \mathcal{L}_{\delta_j}$$

(an \mathcal{O}_U^\times -torsor on U), where

$$\delta_j \stackrel{\text{def}}{=} (\tilde{\Phi}_U^{\log})^{-1}(\log^\Gamma(t_j))$$

Note that \mathcal{M}_j is *independent* of ζ , i.e., it depends only on the *structure morphism* $U^{\log} \rightarrow S^{\log}$ of U^{\log} . Indeed, this follows from the fact (mentioned above) that $\Gamma_S \rightarrow \Gamma_{S_{\mathbb{F}_p}}$ is an *isomorphism*.

Note that γ_U^{\log} (and hence ζ) is determined by the $(\gamma_U^{\log})^{-1}(t_j) \in \Gamma(U, \mathcal{M}_j)$. On the other hand, since γ_U^{\log} and $\tilde{\Phi}_U^{\log}$ are *equal* modulo p (by the definition of F' !), it follows that $(\gamma_U^{\log})^{-1}(t_j)$ and $(\tilde{\Phi}_U^{\log})^{-1}(t_j)$ coincide modulo p , i.e.,

$$(\gamma_U^{\log})^{-1}(t_j) = u_j \cdot (\tilde{\Phi}_U^{\log})^{-1}(t_j)$$

for some unit u_j on U such that $u_j \equiv 1 \pmod{p}$. Thus, γ_U^{\log} (and hence ζ) is completely determined by the collection of functions

$$z_j = \frac{1}{p}(u_j - 1) \in \Gamma(U, \mathcal{O}_U)$$

(Note that it makes sense to speak of “ $\frac{1}{p}$ ” because we are operating under the assumption that U is \mathbf{Z}_p -flat.)

Also, it is easy to see that as ζ varies, all elements of $\Gamma(U, \mathcal{O}_U)$ can occur as z_j . This completes the verification that over this (localized) S^{\log} , F' can indeed be represented over S by d -dimensional affine space (i.e., corresponding to the coordinates z_1, \dots, z_d). This proves (by étale descent) the representability statement of Lemma 1.9; the description of F' modulo p follows immediately from the definitions. \circ

Note that the connection $\nabla_{\mathcal{F}}$ on \mathcal{F} defines a *connection* ∇_F on F which *respects the subfunctor* F' (since it respects the ideal $\Phi_{S_{\mathbf{F}_p}}^* \mathcal{I} \subseteq \Phi_{S_{\mathbf{F}_p}}^* \mathcal{A}$). Thus, we obtain a *connection*

$$\nabla_{F'}$$

on F' . That is to say, F' is a p -adic smooth affine scheme of relative dimension d over S , and $(F', \nabla_{F'})$ forms a *crystal* on $\text{Crys}(S_{\mathbf{F}_p}^{\log})$ (or $\text{Crys}(S^{\log})$) *valued in the category of p -adic affine schemes over the base*. Since $F' \rightarrow S$ is *affine*, we can write

$$(F', \nabla_{F'}) = \text{Spf}(\mathcal{F}', \nabla_{\mathcal{F}'})$$

for some p -adic quasi-coherent \mathcal{O}_S -algebra \mathcal{F}' (with connection $\nabla_{\mathcal{F}'}$). Moreover, we have a natural horizontal morphism $\mathcal{F} \rightarrow \mathcal{F}'$ (derived from $F' \subseteq F$).

Definition 1.10. We shall refer to $(\mathcal{F}', \nabla_{\mathcal{F}'})$ as the *renormalized Frobenius pull-back of the mantle* $(\mathcal{A}, \nabla_{\mathcal{A}})$.

The *reason for the name* is as follows. Suppose that

$$\mathcal{G} \subseteq \mathcal{A}$$

is a *vector bundle* on S of rank $1 + d$ which is *stabilized* by $\nabla_{\mathcal{A}}$, and *contains* $\iota(\mathcal{O}_S)$. Let $\nabla_{\mathcal{G}}$ denote the connection on \mathcal{G} induced by $\nabla_{\mathcal{A}}$, and denote by $\iota_{\mathcal{G}} : \mathcal{O}_S \hookrightarrow \mathcal{G}$ the horizontal morphism induced by $\iota : \mathcal{O}_S \rightarrow \mathcal{A}$. Suppose, moreover, that

$$\mathcal{G} \bigcap \mathcal{I} \rightarrow \mathcal{I}/\mathcal{I}^2$$

is an *isomorphism*.

Definition 1.11. We shall call such a $\mathcal{G} \subseteq \mathcal{A}$ a *geometrizing subbundle*.

Thus, we have an exact sequence (compatible with the connections):

$$0 \rightarrow \mathcal{O}_S \rightarrow \mathcal{G} \rightarrow \mathcal{N} \rightarrow 0$$

(where $(\mathcal{N}, \nabla_{\mathcal{N}})$ is defined so as to make the sequence exact). Note that \mathcal{N} is a rank d *vector bundle* on S . Thus, the *splittings* of this exact sequence define an \mathcal{N}^\vee -*torsor* on S . The corresponding geometric torsor is then Spf of some \mathcal{O}_S -*algebra*

$$\mathrm{Aff}(\mathcal{G})$$

(with $\mathcal{G} \subseteq \mathrm{Aff}(\mathcal{G})$) which we shall call the *affinization* of \mathcal{G} (cf. Definition 1.7). Here, $\mathrm{Aff}(\mathcal{G})$ is locally isomorphic to (the p -adic completion of) a *polynomial algebra* over \mathcal{O}_S in d variables; \mathcal{G} (topologically) *generates* $\mathrm{Aff}(\mathcal{G})$ as a (topological) \mathcal{O}_S -*algebra*; and \mathcal{G} may be recovered from $\mathrm{Aff}(\mathcal{G})$ as the subsheaf of *polynomials of degree ≤ 1* .

Let us define a *Hodge filtration* on \mathcal{G} by:

$$F^\bullet(\mathcal{G}) \stackrel{\mathrm{def}}{=} \mathcal{G} \bigcap F^\bullet(\mathcal{A})$$

Then since $\mathcal{G} = F^1(\mathcal{G}) \oplus \iota_{\mathcal{G}}(\mathcal{O}_S)$, we see that the Hodge filtration defines a (non-horizontal) section of the torsor discussed in the preceding paragraph, which we shall refer to as the *Hodge section*. Thus, the inclusion $\mathcal{G} \hookrightarrow \mathcal{A}$ defines a natural *isomorphism*

$$\mathcal{A} \cong \text{the PD envelope of } \mathrm{Aff}(\mathcal{G}) \text{ at the Hodge section}$$

Next, let us look at *naive Frobenius pull-backs*: We have

$$\Phi_{S_{\mathbb{F}_p}}^* \mathcal{G} \subseteq \Phi_{S_{\mathbb{F}_p}}^* \mathcal{A} = \mathcal{F}$$

Moreover, this inclusion induces a horizontal isomorphism of \mathcal{F} with the PD-envelope of $\mathrm{Aff}(\Phi_{S_{\mathbb{F}_p}}^* \mathcal{G})$ at the pull-back of the Hodge section modulo p .

Finally, we come to *renormalized Frobenius pull-backs*: Let us denote by

$$\mathbf{F}^*(\mathcal{G})$$

the subsheaf of $\Phi_{S_{\mathbf{F}_p}}^*(\mathcal{G}) \otimes_{\mathbf{Z}_p} (\frac{1}{p} \cdot \mathbf{Z}_p)$ consisting of sections whose reduction modulo p is contained in $F^1(-)$. Thus, pulling-back $\iota_{\mathcal{G}}$ by Frobenius, we obtain a horizontal morphism

$$\mathbf{F}^*(\iota_{\mathcal{G}}) : \mathcal{O}_S \hookrightarrow \mathbf{F}^*(\mathcal{G})$$

Just as above, we may then form the *affinization*

$$\mathrm{Aff}(\mathbf{F}^*(\mathcal{G}))$$

Moreover, from $\Phi_{S_{\mathbf{F}_p}}^* \mathcal{G} \hookrightarrow \mathbf{F}^*(\mathcal{G})$, we see that $\mathrm{Aff}(\mathbf{F}^*(\mathcal{G}))$ is naturally an algebra over $\mathrm{Aff}(\Phi_{S_{\mathbf{F}_p}}^* \mathcal{G})$, as well as over \mathcal{F} (since the ideal $p \cdot \mathrm{Aff}(\mathbf{F}^*(\mathcal{G}))$ in $\mathrm{Aff}(\mathbf{F}^*(\mathcal{G}))$ automatically admits divided powers).

Lemma 1.12. *There is a natural horizontal isomorphism*

$$\mathrm{Aff}(\mathbf{F}^*(\mathcal{G})) \cong \mathcal{F}'$$

of \mathcal{F} -algebras over S .

Proof. The proof is similar to that of Lemma 1.9. Indeed, looking back at the proof of Lemma 1.9 shows that essentially what was done there to construct \mathcal{F}' was to adjoin to \mathcal{F} the sections of $\frac{1}{p} \cdot F^1(\mathcal{G})$; but this is exactly the same as the definition of $\mathrm{Aff}(\mathbf{F}^*(\mathcal{G}))$. \square

This shows why we refer to \mathcal{F}' as the *renormalized Frobenius pull-back* of \mathcal{A} .

§1.5. Hodge Subspaces

Now let us return to considering *spiked Frobenius liftings* (Definition 1.1). Recall the subbundle

$$\mathrm{Geo}^{\mathrm{st}}(\mathcal{A}) \subseteq \mathcal{A}$$

of §1.3. Let us endow $\mathrm{Geo}^{\mathrm{st}}(\mathcal{A})$ with the *connection* and *Hodge filtration* induced by those of \mathcal{A} . Write

$$\mathrm{Geo}^{\mathrm{st}}(\mathcal{F}') \stackrel{\mathrm{def}}{=} \mathbf{F}^*(\mathrm{Geo}^{\mathrm{st}}(\mathcal{A}))$$

for the subsheaf of $\Phi_{S_{\mathbf{F}_p}}^*(\mathrm{Geo}^{\mathrm{st}}(\mathcal{A})) \otimes_{\mathbf{Z}_p} (\frac{1}{p} \cdot \mathbf{Z}_p)$ consisting of sections whose reduction modulo p is contained in $F^1(-)$. If we compose the inclusion $\Phi_{S_{\mathbf{F}_p}}^* \mathrm{Geo}^{\mathrm{st}}(\mathcal{A}) \subseteq \Phi_{S_{\mathbf{F}_p}}^* \mathcal{A} = \mathcal{F}$ with $\mathcal{F} \rightarrow \mathcal{F}'$, then one sees (cf. Lemma 1.12) that

$$\Phi_{S_{\mathbf{F}_p}}^* \mathrm{Geo}^{\mathrm{st}}(\mathcal{A}) \rightarrow \mathcal{F}'$$

factors through $\mathrm{Geo}^{\mathrm{st}}(\mathcal{F}') \stackrel{\mathrm{def}}{=} \mathbf{F}^*(\mathrm{Geo}^{\mathrm{st}}(\mathcal{A}))$. Thus, we obtain an inclusion

$$\mathrm{Geo}^{\mathrm{st}}(\mathcal{F}') \hookrightarrow \mathcal{F}'$$

Since the canonical copy of \mathcal{O}_S in \mathcal{F}' is clearly contained inside $\mathrm{Geo}^{\mathrm{st}}(\mathcal{F}')$, it follows that we may construct:

$$(\mathcal{F}')^{\mathrm{st}} \stackrel{\mathrm{def}}{=} \mathrm{Aff}(\mathrm{Geo}^{\mathrm{st}}(\mathcal{F}'))$$

Moreover, the inclusion $\mathrm{Geo}^{\mathrm{st}}(\mathcal{F}') \hookrightarrow \mathcal{F}'$ induces an inclusion

$$(\mathcal{F}')^{\mathrm{st}} \hookrightarrow \mathcal{F}'$$

Thus,

$$(F')^{\mathrm{st}} \stackrel{\mathrm{def}}{=} \mathrm{Spf}((\mathcal{F}')^{\mathrm{st}}) \rightarrow S$$

is *smooth* and *affine* of relative dimension c , and we also have a morphism $F' \rightarrow (F')^{\mathrm{st}}$ which is *smooth* and *affine* of relative dimension $d - c$.

Next, let

$$S'_{\mathbf{F}_p} \rightarrow S_{\mathbf{F}_p}$$

be the *finite inseparable covering* of $S_{\mathbf{F}_p}$ defined by the Frobenius morphism $\Phi_{S_{\mathbf{F}_p}}$ (i.e., $S'_{\mathbf{F}_p}$ is “abstractly” the same as $S_{\mathbf{F}_p}$). Note that the lifting $\Phi : S \rightarrow S$ of $\Phi_{S_{\mathbf{F}_p}}^2$ defines (cf. Lemma 1.9) a morphism

$$\xi : S'_{\mathbf{F}_p} \rightarrow F'_{\mathbf{F}_p} (\subseteq F')$$

whose composite with the structure morphism $F'_{\mathbf{F}_p} \rightarrow S_{\mathbf{F}_p}$ is the inseparable covering $S'_{\mathbf{F}_p} \rightarrow S_{\mathbf{F}_p}$. Now we are ready to make the following:

Definition 1.13. Let $H \subseteq F'$ be a p -adic subscheme satisfying the following conditions:

- (1) the schematic image of ξ in $F'_{\mathbf{F}_p}$ is a k -smooth subscheme $H_{\mathbf{F}_p}^\xi \subseteq H_{\mathbf{F}_p}$;
- (2) the map $H \subseteq F' \rightarrow (F')^{\text{st}}$ is *étale*;
- (3) the composite of the *Kodaira-Spencer morphism* of H (with respect to the connection $\nabla_{F'}$)

$$\mathcal{N}_{H/F'} \rightarrow \Omega_{S^{\log}}$$

(where $\mathcal{N}_{H/F'}$ is the conormal bundle of H in F') with the projection $\Omega_{S^{\log}} \rightarrow \Omega_{S^{\log}}^{\text{wk}}$ is an *isomorphism* $\mathcal{N}_{H/F'} \cong \Omega_{S^{\log}}^{\text{wk}}$.

Then we shall call H a *Hodge subspace for the very ordinary spiked Frobenius lifting* Φ^{\log} .

Note that it follows that H is necessarily \mathbf{Z}_p -flat (since it is *étale* over $(F')^{\text{st}}$, which is S -smooth, hence \mathbf{Z}_p -flat). Moreover, $H_{\mathbf{F}_p}^\xi$ is an \mathbf{F}_p -scheme of dimension d that “lies between” $S'_{\mathbf{F}_p}$ and $S_{\mathbf{F}_p}$; thus, $H_{\mathbf{F}_p}^\xi$ is *finite and purely inseparable* over $S_{\mathbf{F}_p}$. Relative to the discussion of Chapter VII, §3.2, 3.3 (where canonical *modular* spiked Frobenius liftings are discussed), it turns out that $H_{\mathbf{F}_p}^\xi$ will be the scheme “ T ” of *loc. cit.*

Let $H \subseteq F'$ be a *Hodge subspace* for Φ^{\log} . We would like to construct the *renormalized Frobenius pull-back* of F' with respect to H . The construction is entirely similar to that of Lemma 1.9. We start by defining the appropriate *functor*. First note that

$$\Phi_{S_{\mathbf{F}_p}}^*(F', \nabla_{F'})$$

forms a *crystal* valued in the category of p -adic schemes over $\text{Crys}(S_{\mathbf{F}_p}^{\log})$. Alternatively, we may think of $\Phi_{S_{\mathbf{F}_p}}^* F'$ as a p -adic scheme over S equipped with a connection $\Phi_{S_{\mathbf{F}_p}}^* \nabla_{F'}$. As in the context of Lemma 1.9, we would like to regard $\Phi_{S_{\mathbf{F}_p}}^* F'$ as a *functor* on $(\mathbf{Z}_p\text{-flat})$ p -adic log schemes U^{\log} over S^{\log} . Let

$$F'' \subseteq \Phi_{S_{\mathbf{F}_p}}^* F'$$

be the *subfunctor* consisting of those points $(\Phi_{S_{\mathbf{F}_p}}^* F')(U^{\log}) \stackrel{\text{def}}{=} (\Phi_{S_{\mathbf{F}_p}}^* F')(U)$ whose *reduction modulo p* lands inside

$$\Phi_{S_{\mathbf{F}_p}}^* H_{\mathbf{F}_p} \subseteq \Phi_{S_{\mathbf{F}_p}}^* F'$$

Lemma 1.14. *The S -functor F'' is representable by a smooth (hence, \mathbf{Z}_p -flat) p -adic scheme of relative dimension $d = \dim_A(S)$ over S which is an affine space over the base space $\Phi_{S_{\mathbf{F}_p}}^* H$ (where “ $\Phi_{S_{\mathbf{F}_p}}^* H$ ” denotes the pull-back by $\Phi_{S_{\mathbf{F}_p}}$ of the crystal defined by H equipped with the connection on $H \rightarrow S$ induced via the étale morphism $H \rightarrow (F')^{\text{st}}$ from the connection on $(F')^{\text{st}} \rightarrow S$).*

Proof. The proof is entirely similar to that of Lemma 1.9. Whereas in the case of Lemma 1.9, we adjoined $\frac{1}{p}$ times generators of the diagonal ideal, in the present case, we adjoin $\frac{1}{p}$ times generators of the ideal defining H . Since (cf. Definition 1.13, (2)), H may be thought of as corresponding to the *strong* dimensions of F' , it thus follows that F'' is obtained by adjoining $\frac{1}{p}$ times the *weak variables* (to $\Phi_{S_{\mathbf{F}_p}}^* \mathcal{F}'$). \circ

Note that since $\Phi_{S_{\mathbf{F}_p}}^* H_{\mathbf{F}_p}$ is a *horizontal* subscheme of $\Phi_{S_{\mathbf{F}_p}}^* F'$, it follows that F'' admits a *connection*

$$\nabla_{F''}$$

induced from that of $\Phi_{S_{\mathbf{F}_p}}^* F'$. By Condition (1) of Definition 1.13, the image of ξ lands inside $H_{\mathbf{F}_p}$. Thus, by the *definition* of the subfunctor $F'' \subseteq \Phi_{S_{\mathbf{F}_p}}^* F'$, the Frobenius lifting $\Phi^{\log} : S^{\log} \rightarrow S^{\log}$ defines a section

$$\sigma_{\Phi} : S \rightarrow F''$$

Let F''_{Φ} be the *PD-envelope* of F'' at the image of σ_{Φ} . Then $\nabla_{F''}$ extends to a connection $\nabla_{F''_{\Phi}}$ on F''_{Φ} . We shall sometimes (by abuse of notation) regard σ_{Φ} as a morphism $S \rightarrow F''_{\Phi}$.

Note that the *Kodaira-Spencer morphism*

$$\sigma_{\Phi}^* \Omega_{F''/S} \rightarrow \Omega_{S^{\log}}$$

of σ_{Φ} (obtained by differentiating σ_{Φ} with respect to the connection $\nabla_{F''}$) is an *isomorphism*. Indeed, after unraveling all the definitions, this Kodaira-Spencer morphism essentially amounts to $d\Phi^{\log}$, except that we divided (all of) $\Omega_{S^{\log}}$ once by p (in passing from F to $F' \subseteq F$), and then we divided the “weak variables of $\Omega_{S^{\log}}$ ” (i.e., the variables corresponding to the subsheaf $\Omega_{S_{\mathbf{F}_p}}^{\log}$ – cf. §1.1) once more by p (in passing from $\Phi_{S_{\mathbf{F}_p}}^* F'$ to F''). Thus, the fact that this Kodaira-Spencer morphism is an isomorphism is a consequence of Condition (3) of Definition 1.1.

Lemma 1.15. *There is a horizontal isomorphism*

$$\mathrm{Spf}(\mathcal{A}) \cong F''_{\Phi}$$

that maps $V(\mathcal{I})$ (i.e., the diagonal in $\mathrm{Spf}(\mathcal{A})$) to σ_{Φ} .

Proof. This is a formal consequence of what we have already done. Indeed, to define such an isomorphism, let π_L, π_R be the two projections from $\mathrm{Spf}(\mathcal{A})$ to S^{\log} defined by the left and right inclusions of \mathcal{O}_S into \mathcal{A} , respectively. Then if we compose the section

$$\pi_L^* \sigma_{\Phi} : \mathrm{Spf}(\mathcal{A}) \rightarrow \pi_L^* F''_{\Phi}$$

(obtained by pulling $\sigma_{\Phi} : S \rightarrow F''_{\Phi}$ back via π_L) with

$$\pi_L^* F''_{\Phi} \cong \pi_R^* F''_{\Phi} \rightarrow F''_{\Phi}$$

(where the isomorphism is defined by the connection $\nabla_{F''_{\Phi}}$, and the final morphism is just the projection to F''_{Φ}), we obtain an S -morphism

$$\mathrm{Spf}(\mathcal{A}) \rightarrow F''_{\Phi}$$

that maps $V(\mathcal{I})$ to σ_{Φ} , as desired. In fact, this morphism is just the morphism that maps the \mathcal{O}_S on the left of the “ \mathcal{A} inside \mathcal{F}''_{Φ} ” (i.e., the copy of \mathcal{A} that we pulled back twice by Frobenius to form F''_{Φ}) to the \mathcal{O}_S on the left of the “present \mathcal{A} ” (i.e., of $\mathrm{Spf}(\mathcal{A}) \rightarrow F''_{\Phi}$) by $\Phi^{-1} : \mathcal{O}_S \rightarrow \mathcal{O}_S$.

The fact that this morphism is *horizontal* follows from the integrability of $\nabla_{F''_{\Phi}}$ (and the fact that, with our conventions, pulling back on the left is always a “horizontal operation”). That it is an *isomorphism* follows from the fact that the Kodaira-Spencer morphism of σ_{Φ} (which, if one sorts through the definitions, is simply the morphism between the respective cotangent bundles of $V(\mathcal{I})$ and σ_{Φ}) is an isomorphism.

○

We are now ready to define a “formal \mathcal{MF}^{∇} -object” in the spirit of [Falt1], §2. Unlike [Falt1], §2, *this object will not be of finite rank*, so we cannot immediately apply the theory of [Falt1], §2, to obtain a *Galois representation*. That is to say, we must resort to various tricks – cf. the technique of *crystalline induction* discussed in [Mzk1], Chapter V – to define such a representation – cf. §2.3.

We work over S^{\log} , and assume that we have been given a *Hodge subspace* H for Φ^{\log} . Then we let

$$\Xi \stackrel{\text{def}}{=} \mathrm{Spf}(\mathcal{A}) \times_S F'$$

With its natural *connection* ∇_Ξ (obtained by taking the product of $\nabla_{\mathcal{A}}$ and $\nabla_{F'}$), (Ξ, ∇_Ξ) forms a *crystal* on $\mathrm{Crys}(S^{\mathrm{log}})$. The *Hodge subspace* H_Ξ of Ξ is defined to be

$$H_\Xi \stackrel{\text{def}}{=} H = S \times_S H \subseteq \Xi$$

(where the “ S ” on the left is the *diagonal* $V(\mathcal{I})$ inside $\mathrm{Spf}(\mathcal{A})$). By taking the *renormalized Frobenius pull-back* $\mathbf{F}^*(\Xi, \nabla_\Xi)$ with respect to this subspace (i.e., the “product” of what was done in Lemmas 1.9 and 1.14), we thus obtain

$$F' \times_S F'' \longleftarrow F' \times_S F''_\Phi \cong \mathrm{Spf}(\mathcal{A}) \times_S F' = \Xi$$

where the first morphism is the natural PD-completion morphism, and the isomorphism is that obtained by applying Lemma 1.15 and switching the two factors. We shall abbreviate these last steps and simply write:

$$\Xi_\Phi : \mathbf{F}^*(\Xi, \nabla_\Xi) \cong (\Xi, \nabla_\Xi)$$

We state this as a Theorem:

Theorem 1.16. *Let $\Phi^{\mathrm{log}} : S^{\mathrm{log}} \rightarrow S^{\mathrm{log}}$ be a very ordinary spiked Frobenius lifting (Definition 1.1), and let $H \subseteq \mathbf{F}^*(\mathcal{A}, \nabla_{\mathcal{A}})$ be a p -adic subscheme of the renormalized Frobenius pull-back of the mantle of S^{log} which forms a Hodge subspace for Φ^{log} (cf. Definition 1.13). Then there is a crystal*

$$(\Xi, \nabla_\Xi)$$

(on $\mathrm{Crys}(S^{\mathrm{log}})$) in p -adic schemes, together with a Hodge subspace $H_\Xi \subseteq \Xi$ with respect to which there is a horizontal isomorphism

$$\Xi_\Phi : \mathbf{F}^*(\Xi, \nabla_\Xi) \cong (\Xi, \nabla_\Xi)$$

We shall refer to this data $(\Xi, \nabla_\Xi, H_\Xi, \Xi_\Phi)$ as the formal uniformizing \mathcal{MF}^∇ -object associated to (Φ^{log}, H) .

§2. Associated Galois Representations

We maintain the notation of the preceding §. In particular, we assume that we have been given a *very ordinary spiked Frobenius lifting* $\Phi : S^{\log} \rightarrow S^{\log}$ (cf. Definition 1.1), together with a *Hodge subspace* $H \subseteq F'$ (cf. Definition 1.13).

§2.1. The Strictly Weak Pair of Frobenius Liftings over the Strong Perfection

Recall the projective system discussed at the end of §1.2:

$$\dots \rightarrow S^{\text{st}}[n+1] \rightarrow S^{\text{st}}[n] \rightarrow \dots \rightarrow S^{\text{st}}[0] = S$$

Let us write

$$S^{\text{st}} \rightarrow S$$

for the *the projective limit* of this system, and

$$\widehat{S}^{\text{st}} \rightarrow S$$

for the *p-adic completion* of S^{st} . Let $(S^{\text{st}})^{\log}$ (respectively, $(\widehat{S}^{\text{st}})^{\log}$) be the result of equipping S^{st} (respectively, \widehat{S}^{st}) with the log structure pulled back from S . Thus, we have a morphism

$$\epsilon^{\log} : (\widehat{S}^{\text{st}})^{\log} \rightarrow S^{\log}$$

Moreover, because of the relationship between G_{Φ}^{st} and Φ^{\log} (as in Theorem 1.5), one sees easily that Φ^{\log} naturally induces a morphism

$$\Phi_{\epsilon}^{\log} : (\widehat{S}^{\text{st}})^{\log} \rightarrow (\widehat{S}^{\text{st}})^{\log}$$

such that $\epsilon^{\log} \circ \Phi_{\epsilon}^{\log} = \Phi^{\log} \circ \epsilon^{\log}$.

Definition 2.1. We shall call $(S^{\text{st}})^{\log}$ (respectively, $(\widehat{S}^{\text{st}})^{\log}$) the *strong perfection* (respectively, *completed strong perfection*) of S^{\log} (with respect to Φ^{\log}).

Now let us see what happens when we *pull things back* by ϵ . We start with $\epsilon^* \text{Geo}^{\text{st}}(\mathcal{A})$ (cf. §1.3). The key point now is to observe that *the Hodge filtration*

$$\epsilon^* F^1(\mathrm{Geo}^{\mathrm{st}}(\mathcal{A})) \subseteq \epsilon^* \mathrm{Geo}^{\mathrm{st}}(\mathcal{A})$$

is now *horizontal*. (Indeed, one can see this by, for instance, localizing at a point $z \in U(k)$ and using Theorem 1.5 to reduce the question to a question on \mathcal{G}_2 , where the result is clear – cf. the construction of a splitting of Dieudonné crystals in the discussion following Chapter VIII, Definition 2.10.) Let

$$\mathcal{A}_\epsilon \stackrel{\mathrm{def}}{=} \epsilon^* \mathcal{A} / \{\text{the PD ideal generated by } \epsilon^* F^1(\mathrm{Geo}^{\mathrm{st}}(\mathcal{A}))\}$$

Then \mathcal{A}_ϵ inherits from $\epsilon^* \mathcal{A}$ a *connection*, together with a *Hodge filtration* induced by the section

$$\sigma_{\mathcal{A}} : \widehat{S}^{\mathrm{st}} \rightarrow \mathrm{Spf}(\mathcal{A}_\epsilon)$$

defined by the *diagonal* $\epsilon^* V(\mathcal{I}) \subseteq \mathrm{Spf}(\epsilon^* \mathcal{A})$ (which in fact lies in $\mathrm{Spf}(\mathcal{A}_\epsilon) \subseteq \mathrm{Spf}(\epsilon^* \mathcal{A})$, since $F^1(\mathrm{Geo}^{\mathrm{st}}(\mathcal{A})) \subseteq F^1(\mathcal{A}) = \mathcal{I}$).

Now since $\epsilon^* F^1(\mathrm{Geo}^{\mathrm{st}}(\mathcal{A})) \subseteq \epsilon^* \mathrm{Geo}^{\mathrm{st}}(\mathcal{A})$ is *horizontal*, it defines (after multiplication by p^{-1}) an *ideal* in $\epsilon^*(\mathcal{F}')^{\mathrm{st}}$ such that the quotient of $\epsilon^*(\mathcal{F}')^{\mathrm{st}}$ by this ideal gives rise to a *horizontal section*

$$\sigma_{\mathcal{F}'}^{\mathrm{st}} : \widehat{S}^{\mathrm{st}} \rightarrow \epsilon^*(F')^{\mathrm{st}}$$

over $\widehat{S}^{\mathrm{st}}$. Moreover, it follows from the fact that the Frobenius lifting Φ^{\log} preserves the Hodge filtration $F^1(P^{\mathrm{st}}) \subseteq P^{\mathrm{st}}$ (cf. §1.2) that if $\xi : S'_{\mathbf{F}_p} \rightarrow F'_{\mathbf{F}_p}$ is as in §1.5, then $\epsilon^* \xi$ maps into the image of $\sigma_{\mathcal{F}'}^{\mathrm{st}}$. Let

$$F'_\epsilon \stackrel{\mathrm{def}}{=} \mathrm{Spf}(\mathcal{F}'_\epsilon) \stackrel{\mathrm{def}}{=} (\epsilon^* F') \times_{\epsilon^*(F')^{\mathrm{st}}, \sigma_{\mathcal{F}'}^{\mathrm{st}}} \widehat{S}^{\mathrm{st}} \subseteq \epsilon^* F'$$

Note that (since $\sigma_{\mathcal{F}'}^{\mathrm{st}}$ is *horizontal*) \mathcal{F}'_ϵ inherits a *connection* from $\epsilon^* \mathcal{F}'$. Next, let us observe that

$$\epsilon^* H \rightarrow \epsilon^*(F')^{\mathrm{st}}$$

is *étale* (by Definition 1.13, (2)). Thus,

$$H_\epsilon \stackrel{\mathrm{def}}{=} (\epsilon^* H) \times_{\epsilon^*(F')^{\mathrm{st}}, \sigma_{\mathcal{F}'}^{\mathrm{st}}} \widehat{S}^{\mathrm{st}} \subseteq (\epsilon^* F') \times_{\epsilon^*(F')^{\mathrm{st}}, \sigma_{\mathcal{F}'}^{\mathrm{st}}} \widehat{S}^{\mathrm{st}} = F'_\epsilon$$

is *étale* over $\widehat{S}^{\mathrm{st}}$. Moreover, since $\epsilon^* \xi$ maps into the image of $\sigma_{\mathcal{F}'}^{\mathrm{st}}$ (as well as into $\epsilon^* H$ – by Definition 1.13, (1)), it follows that $\epsilon^* \xi$ maps into H_ϵ . Thus, H_ϵ contains a connected component which is *finite and purely*

inseparable over \widehat{S}^{st} . Since it is also *étale* over \widehat{S}^{st} , it follows that this connected component maps *isomorphically* to \widehat{S}^{st} , hence defines a *section*

$$\sigma_{\mathcal{F}} : \widehat{S}^{\text{st}} \rightarrow F'_\epsilon$$

This section $\sigma_{\mathcal{F}}$ then induces a *Hodge filtration* on \mathcal{F}'_ϵ .

Finally, observe that \mathcal{F}'_ϵ (respectively, \mathcal{A}_ϵ) is (non-canonically) *isomorphic* to the p -adic completion of (respectively, the p -adic completion of the PD-envelope at the origin of) the *polynomial algebra* $\mathcal{O}_S[t_1, \dots, t_{d-c}]$ (where the t_j are indeterminates).

Lemma 2.2. *There is a natural isomorphism*

$$\Omega_{(\widehat{S}^{\text{st}})^{\log}} \cong \epsilon^* \Omega_{S^{\log}}^{\text{wk}}$$

where $\Omega_{S^{\log}}^{\text{wk}} \stackrel{\text{def}}{=} \Omega_{S^{\log}} / \Omega_{S^{\log}}^{\text{st}}$. (Here, our differentials are, as always, p -adically continuous.)

Proof. At any rate, $d\epsilon^{\log}$ induces a morphism $\epsilon^* \Omega_{S^{\log}} \rightarrow \Omega_{(\widehat{S}^{\text{st}})^{\log}}$ which clearly vanishes on $\epsilon^* \Omega_{S^{\log}}^{\text{st}}$. The fact that this gives rise to an isomorphism as claimed follows, for instance, by computing locally using Theorem 1.5. \circ

Next, we would like to consider the *Kodaira-Spencer morphisms* of the *Hodge sections* $\sigma_{\mathcal{A}}$ and $\sigma_{\mathcal{F}}$ constructed above. The Kodaira-Spencer morphisms are given by morphisms

$$\mathcal{N}_{\sigma_{\mathcal{A}}} \rightarrow \Omega_{(\widehat{S}^{\text{st}})^{\log}}; \quad \mathcal{N}_{\sigma_{\mathcal{F}}} \rightarrow \Omega_{(\widehat{S}^{\text{st}})^{\log}}$$

(where $\mathcal{N}_{\sigma_{\mathcal{A}}}$ (respectively, $\mathcal{N}_{\sigma_{\mathcal{F}}}$) is the conormal bundle of the embedding $\sigma_{\mathcal{A}}$ (respectively, $\sigma_{\mathcal{F}}$)). I *claim* that these Kodaira-Spencer morphisms are *isomorphisms*. Indeed, for the first morphism, this follows immediately from the fact that the Kodaira-Spencer morphism for the diagonal $V(\mathcal{I}) \subseteq \text{Spf}(\mathcal{A})$ “is” the identity morphism, hence an isomorphism. For the second morphism, this follows from Lemma 2.2 above and Condition (3) of Definition 1.13. This completes the proof of the claim.

Now observe that from Theorem 1.16, we have *isomorphisms*

$$\mathbf{F}^*(\mathcal{A}_\epsilon) \cong \mathcal{F}'_\epsilon; \quad \mathcal{A}_\epsilon \cong \{\text{some PD envelope of } \mathbf{F}^*(\mathcal{F}'_\epsilon)\}$$

Thus, by using the fact that the Kodaira-Spencer morphisms (discussed in the preceding paragraph) are isomorphisms, it follows by a standard deformation argument that there exists a unique pair of *Frobenius liftings*

$$\Phi_0^{\log}, \Phi_1^{\log} : (\widehat{S}^{\text{st}})^{\log} \rightarrow (\widehat{S}^{\text{st}})^{\log}$$

such that the morphism $\Phi_0^* \mathcal{A}_\epsilon \rightarrow \mathcal{F}'_\epsilon$ (respectively, $\Phi_1^* \mathcal{F}'_\epsilon \rightarrow \mathcal{A}_\epsilon$) induced by the Frobenius action of Theorem 1.16 *respects the ideals defining the Hodge sections*. Indeed, F'_ϵ itself may be regarded as a sort of *moduli space of Frobenius liftings* (cf. Lemma 1.9). Thus, the section $\sigma_{\mathcal{F}} : \widehat{S}^{\text{st}} \rightarrow F'_\epsilon$ corresponds (via this modular interpretation of F'_ϵ) to a Frobenius lifting Φ_0^{\log} of the sort desired. Similarly, since the Kodaira-Spencer morphism of $\sigma_{\mathcal{F}}$ is an *isomorphism*, the renormalized Frobenius pull-back of F'_ϵ (relative to the Hodge section $\sigma_{\mathcal{F}}$) is again a sort of *moduli space of Frobenius liftings* (cf. the proof of Lemma 1.9). Thus, the section $\sigma_{\mathcal{A}}$ defines (via the above isomorphism of \mathcal{A}_ϵ with the renormalized Frobenius pull-back of \mathcal{F}'_ϵ) a Frobenius lifting Φ_1^{\log} of the sort desired.

Next, observe that the fact that the Kodaira-Spencer morphisms discussed above are isomorphisms implies that

$$\Omega_{\Phi_0^{\log}} \stackrel{\text{def}}{=} \frac{1}{p} \cdot d\Phi_0^{\log}; \quad \Omega_{\Phi_1^{\log}} \stackrel{\text{def}}{=} \frac{1}{p} \cdot d\Phi_1^{\log}$$

are *isomorphisms*. Put another way, the pair of Frobenius liftings

$$\Phi_0^{\log}, \Phi_1^{\log} : (\widehat{S}^{\text{st}})^{\log} \rightarrow (\widehat{S}^{\text{st}})^{\log}$$

effectively form a *binary-ordinary anabelian Frobenius system of length 2 and order 1 on $(\widehat{S}^{\text{st}})^{\log}$* . Indeed, the only difference between our situation here and that studied extensively in Chapter VIII, §2, is that $(\widehat{S}^{\text{st}})^{\log}$ is not formally smooth of finite type over A . However, it is easy to see that this difference is cosmetic: Indeed, either one can go back through Chapter VIII, §2 and redevelop the theory there for the present case over a log scheme like $(\widehat{S}^{\text{st}})^{\log}$ (the proofs and constructions go through exactly as before). Or, one can localize at a point $z \in U(k)$, and (in the language of Theorem 1.5) restrict to a fiber of Γ_z^{st} over the A -valued point of Θ_z^{st} given by the identity; such a fiber will then be (the completion at z of) a log scheme which is formally smooth and of finite type over A , so one can (almost) literally apply the theory of Chapter VIII, §2, to the restrictions of Φ_1^{\log} and Φ_0^{\log} to this fiber. At any rate, we state this as a Theorem:

Theorem 2.3. *Associated to (Φ^{\log}, H) , we have a natural pair of Frobenius liftings*

$$\Phi_0^{\log}, \Phi_1^{\log} : (\widehat{S}^{\text{st}})^{\log} \rightarrow (\widehat{S}^{\text{st}})^{\log}$$

on $(\widehat{S}^{\text{st}})^{\log}$ which forms a binary-ordinary anabelian Frobenius system of length 2 and order 1 on $(\widehat{S}^{\text{st}})^{\log}$. In particular, one can thus apply all the constructions (formal group-theoretic and affine multi-uniformizations, multi-canonical points, etc.) discussed in Chapter VIII to this pair over $(\widehat{S}^{\text{st}})^{\log}$. We shall call this pair of Frobenius liftings the strictly weak pair of Frobenius liftings on $(\widehat{S}^{\text{st}})^{\log}$.

Finally, the composite $\Phi_0^{\log} \circ \Phi_1^{\log} : (\widehat{S}^{\text{st}})^{\log} \rightarrow (\widehat{S}^{\text{st}})^{\log}$ is equal to the morphism $\Phi_{\epsilon}^{\log} : (\widehat{S}^{\text{st}})^{\log} \rightarrow (\widehat{S}^{\text{st}})^{\log}$ naturally induced by Φ^{\log} .

Proof. It remains only to verify the final statement concerning the composite. But this follows from the fact that both the composite and Φ_{ϵ}^{\log} preserve the Hodge section σ_A of $\text{Spf}(\mathcal{A}_{\epsilon})$. Thus, if one computes the effect of the pull-back morphism (induced by the Frobenius action of Theorem 1.16)

$$\mathcal{A}_{\epsilon} \rightarrow (\Phi_{\widehat{S}_{\mathbf{F}_p}^{\text{st}}}^2)^* \mathcal{A}_{\epsilon} = \Phi_{\epsilon}^* \mathcal{A}_{\epsilon} = \Phi_1^* \Phi_0^* \mathcal{A}_{\epsilon} \rightarrow \mathcal{A}_{\epsilon}$$

on the (horizontal) sections of \mathcal{A}_{ϵ} arising from the inclusion $\mathcal{O}_S \hookrightarrow \mathcal{A}$ from the left, one sees that the morphisms $\Phi_0 \circ \Phi_1$ and Φ_{ϵ} are the same.

○

Before continuing, we note the following consequence of the above discussion:

Lemma 2.4. *The section $\xi : S'_{\mathbf{F}_p} \rightarrow F'_{\mathbf{F}_p}$ (appearing in Definition 1.13) “factors through $T_{\mathbf{F}_p}$ ” (where $T \rightarrow S$ is as defined at the end of §1.2) in the following sense: There exists a unique section*

$$\zeta : T_{\mathbf{F}_p} \rightarrow F'_{\mathbf{F}_p}$$

(whose composite with $F'_{\mathbf{F}_p} \rightarrow S_{\mathbf{F}_p}$ is the natural morphism $T_{\mathbf{F}_p} \rightarrow S_{\mathbf{F}_p}$) such that $\text{Im}(\zeta) = \text{Im}(\xi)$. We shall denote the composite of ζ with $F'_{\mathbf{F}_p} \rightarrow (F')_{\mathbf{F}_p}^{\text{st}}$ by $\zeta^{\text{st}} : T_{\mathbf{F}_p} \rightarrow (F')_{\mathbf{F}_p}^{\text{st}}$.

Proof. Uniqueness follows from the condition that $\text{Im}(\zeta) = \text{Im}(\xi)$, plus the fact both $T_{\mathbf{F}_p}$ and $S'_{\mathbf{F}_p}$ are radicial over $S_{\mathbf{F}_p}$. As for existence, if $H_{\mathbf{F}_p}^{\xi} \stackrel{\text{def}}{=} \text{Im}(\xi)$ is as in Definition 1.13, then we simply observe that by

the discussion preceding Lemma 2.2 (concerning the construction of $\sigma_{\mathcal{F}}$), it follows that $\mathcal{O}_{H_{\mathbb{F}_p}^{\xi}} \subseteq \mathcal{O}_{(\widehat{S}^{\text{st}})_{\mathbb{F}_p}}$. Thus, we have that

$$\mathcal{O}_{S_{\mathbb{F}_p}} \subseteq \mathcal{O}_{H_{\mathbb{F}_p}^{\xi}} \subseteq \mathcal{O}_{(\widehat{S}^{\text{st}})_{\mathbb{F}_p}} \bigcap \mathcal{O}_{S'_{\mathbb{F}_p}} \subseteq \mathcal{O}_{T_{\mathbb{F}_p}}$$

as desired. \circ

Finally, we close this § by making explicit certain *consequences of Theorem 2.3*. Note that, just as in Chapter VIII, §2.2 (cf. especially the discussion preceding Definition 2.9), we can construct from Φ_0^{\log} , Φ_1^{\log} a *canonical weak uniformizing \mathcal{MF}^{∇} -object*

$$(P, \nabla_P, F^*(P), \Phi_P)_{(\widehat{S}^{\text{st}})^{\log}}^{\text{wk}}$$

on $(\widehat{S}^{\text{st}})^{\log}$. This object gives rise to a *canonical weak uniformizing (log) p -divisible group* (cf. Chapter VIII, Definition 2.9) $(G_{\Phi})_{(\widehat{S}^{\text{st}})^{\log}}$ on $(\widehat{S}^{\text{st}})^{\log}$ which fits into an exact sequence

$$0 \rightarrow (G_{\Omega_{S^{\log}}^{\text{wk,et}}})_{(\widehat{S}^{\text{st}})^{\log}} \rightarrow (G_{\Phi})_{(\widehat{S}^{\text{st}})^{\log}} \rightarrow \mathbf{Q}_p/\mathbf{Z}_p \rightarrow 0$$

(where $(G_{\Omega_{S^{\log}}^{\text{wk,et}}})_{(\widehat{S}^{\text{st}})^{\log}}$ is the tensor product over \mathbf{Z}_p of the p -divisible group $\mathbf{G}_m[p^{\infty}]$ (associated to the multiplicative group) with the étale local system $\Omega_{S^{\log}}^{\text{wk,et}}|_{(\widehat{S}^{\text{st}})^{\log}}$).

Note, moreover, that for each $n \geq 1$, the object $(P, \nabla_P, F^*(P), \Phi_P)_{(\widehat{S}^{\text{st}})^{\log}}^{\text{wk}} \otimes \mathbf{Z}/p^{n+1}\mathbf{Z}$ descends to an object

$$(P, \nabla_P, F^*(P),$$

$$\Phi_P)_{(S^{\text{st}})^{\log}}^{\text{wk}} \otimes \mathbf{Z}/p^{n+1}\mathbf{Z}$$

over $(S^{\text{st}})^{\log}$ (i.e., tensoring with $\mathbf{Z}/p^{n+1}\mathbf{Z}$ yields an object which is “already p -adically complete,” hence is defined over $(S^{\text{st}})^{\log}$). Thus, the kernel of multiplication by p^n on the above exact sequence of (log) p -divisible groups gives rise to an *exact sequence of finite, flat (log) group schemes*

$$0 \rightarrow (G_{\Omega_{S^{\log}}^{\text{wk,et}}})_{(S^{\text{st}})^{\log}}[p^n] \rightarrow (G_{\Phi})_{(S^{\text{st}})^{\log}}[p^n] \rightarrow \mathbf{Q}_p/\mathbf{Z}_p[p^n] \rightarrow 0$$

over $(S^{\text{st}})^{\log}$. Let us denote by

$$(T_n^{\text{wk}})^{\log} \rightarrow (S^{\text{st}})^{\log}$$

the finite, flat morphism (of degree $p^{2n \cdot (d-c)}$) of *splittings* of the above exact sequence of finite, flat (log) group schemes. Note that this morphism is *log étale in characteristic zero*. Thus, it corresponds to a morphism of *sets with a continuous action of $\Pi_{S^{\log}} \stackrel{\text{def}}{=} \pi_1(S_K^{\log}, s)$* :

$$\mathcal{V}_n^{\text{wk}} \rightarrow \mathcal{U}^{\text{st}}$$

Let us write

$$\mathcal{V} \stackrel{\text{def}}{=} \varprojlim \mathcal{V}_n^{\text{wk}}$$

(where the limit is over integers $n \geq 1$) – cf. Chapter VIII, Definition 1.4, Theorem 2.12.

Definition 2.5. We shall refer to \mathcal{V} (respectively, \mathcal{U}^{st}) as the *set-theoretic (respectively, set-theoretic strong) canonical Galois representation associated to Φ^{\log} (and H)*.

We close with the following “spiked analogue” of Chapter VIII, Theorem 2.12:

Corollary 2.6. *For every $m \geq 1$, the $\Pi_{S^{\log}}$ -set corresponding to the covering*

$$(\Phi^{\log})^m : S^{\log} \rightarrow S^{\log}$$

is a quotient of the $\Pi_{S^{\log}}$ -set \mathcal{V} .

Proof. This is essentially an immediate corollary of the above discussion, combined with Chapter VIII, Theorem 2.12. Indeed, the only slight technical issue (relative to applying Chapter VIII, Theorem 2.12) that one needs to worry about arises from the fact that the Frobenius liftings Φ_0^{\log} and Φ_1^{\log} are *only defined over the p-adic completion $(\widehat{S}^{\text{st}})^{\log}$ (i.e., not over $(S^{\text{st}})^{\log}$ itself)*. That is to say, one knows that the covering

$$(\Phi_0^{\log} \circ \Phi_1^{\log})_K^m : (\widehat{S}^{\text{st}})_K^{\log} \rightarrow (\widehat{S}^{\text{st}})_K^{\log}$$

is a *quotient* of the covering

$$\{(T_{2m}^{\text{wk}})^{\log} \times_{(S^{\text{st}})^{\log}} (\widehat{S}^{\text{st}})^{\log}\}_K \rightarrow (\widehat{S}^{\text{st}})^{\log}_K$$

by Chapter VIII, Theorem 2.12 (applied to $(\widehat{S}^{\text{st}})^{\log}$), and one would like to *descend* this statement to a statement concerning coverings of $(S^{\text{st}})^{\log}_K$. But the fact that such a descent is always possible is an immediate consequence of “Hensel’s lemma” applied to *solutions* in $\mathcal{O}_{\widehat{S}^{\text{st}}}$ of monic polynomials (whose discriminant divides a power of p) with *coefficients* in $\mathcal{O}_{S^{\text{st}}}$. \bigcirc

§2.2. The Associated Non-affine Geometry

Let $T^{\log} \rightarrow S^{\log}$ be $S^{\text{st}}[1]^{\log} \rightarrow S^{\text{st}}[0]^{\log}$ as in the discussion following Theorem 1.5. Let

$$\mathcal{A}_T$$

be the pull-back of \mathcal{A} to T . We equip \mathcal{A}_T with the *connection*, *Hodge filtration*, and *Frobenius action* pulled back from those of \mathcal{A} . Next, we would like to construct an object \mathcal{B}_T from \mathcal{F}' as follows: First, we pull \mathcal{F}' back to T to obtain \mathcal{F}'_T . Then we let

$$\mathcal{B}_T$$

be the *p-adic completion of the PD-envelope of \mathcal{F}'_T at the image of the section ζ* of Lemma 2.4. Since the PD-envelope \mathcal{B}_T is a more *similar* sort of object to \mathcal{A}_T than the affine space \mathcal{F}'_T , it is often more natural to work with \mathcal{B}_T than with \mathcal{F}'_T . Naturally, the connection on \mathcal{F}' induces a natural *connection* $\nabla_{\mathcal{B}_T}$ on \mathcal{B}_T . Moreover, the *Hodge subspace* $H \subseteq F'$ induces a Hodge subspace

$$H_{\mathcal{B}_T} \subseteq \mathcal{B}_T \stackrel{\text{def}}{=} \text{Spf}(\mathcal{B}_T)$$

which, in turn, induces (by taking divided powers of the PD-ideal $\mathcal{I}_{H_{\mathcal{B}_T}}$ that defines $H_{\mathcal{B}_T}$) a *Hodge filtration* on \mathcal{B}_T . Similarly, we write

$$(\mathcal{B}_{S'}, \nabla_{\mathcal{B}_{S'}})$$

for the *p-adic completion of the PD-envelope of $\mathcal{F}'_{S'}$* $\stackrel{\text{def}}{=} \Phi_{S_{\mathbf{F}_p}}^* \mathcal{F}'$ at the image of the section $\xi : S'_{\mathbf{F}_p} \rightarrow F'_{\mathbf{F}_p}$, and

$$H_{\mathcal{B}_{S'}} \subseteq \mathcal{B}_{S'} \stackrel{\text{def}}{=} \text{Spf}(\mathcal{B}_{S'})$$

for the corresponding *Hodge subspace*. Finally, let us note that by taking the p -adic completion of the PD-envelope (at the images of ζ and ξ , respectively) of the pull-backs of the sub-crystal $(\mathcal{F}')^{\text{st}} \subseteq \mathcal{F}'$ to $T_{\mathbb{F}_p}$ and $S'_{\mathbb{F}_p}$, we obtain *PD-subalgebras*

$$\mathcal{B}_T^{\text{st}} \subseteq \mathcal{B}_T; \quad \mathcal{B}_{S'}^{\text{st}} \subseteq \mathcal{B}_{S'}$$

Next, we would like to discuss various *natural subquotient sheaves* of \mathcal{A} and \mathcal{B}_T . Let

$$\mathcal{E} \stackrel{\text{def}}{=} \mathcal{A}/\mathcal{A}^{\text{st}}$$

Note that since \mathcal{A}^{st} is stabilized by $\nabla_{\mathcal{A}}$ and $\Phi_{\mathcal{A}}$ (cf. §1.3), it follows that \mathcal{E} gets a natural *connection* $\nabla_{\mathcal{E}}$ and *Frobenius action* $\Phi_{\mathcal{E}}$ (induced by $\nabla_{\mathcal{A}}$ and $\Phi_{\mathcal{A}}$, respectively). Note that \mathcal{E} is a *quotient* not only of \mathcal{A} , but also of the sheaf \mathcal{D} considered in §1.3. It thus follows that the Frobenius action $\Phi_{\mathcal{E}}$ is *divisible by p* .

Lemma 2.7. (*Assume that Φ^{\log} admits a Hodge subspace H (cf. Definition 1.13). Then the Frobenius action $\Phi_{\mathcal{E}}$ on \mathcal{E} is divisible by p^2 .*)

Proof. Recall that the isomorphism of Lemma 1.15 defines a morphism from the *renormalized Frobenius pull-back* of \mathcal{F}' to \mathcal{A} . Since there is (by definition) a natural morphism from the naive Frobenius pull-back of \mathcal{F}' to the renormalized Frobenius pull-back of \mathcal{F}' , and, moreover, $\mathcal{B}_{S'}$ is defined as the (completion of the) PD-envelope of this naive pull-back, we thus see that the isomorphism of Lemma 1.15 gives rise to a Frobenius morphism

$$\mathcal{B}_{S'} \rightarrow \mathcal{A}$$

which *factors through* (the PD-envelope at σ_{Φ} of) \mathcal{F}'' , hence is *zero modulo p* on the sections in the ideal

$$(\mathcal{I}_{H_{B_{S'}}})_{\mathbb{F}_p} \subseteq (\mathcal{B}_{S'})_{\mathbb{F}_p}$$

(where $\mathcal{I}_{H_{B_{S'}}} \subseteq \mathcal{B}_{S'}$ is the PD-ideal defining the Hodge subspace $H_{B_{S'}} \subseteq B_{S'}$).

On the other hand, recall from the proof of Lemma 1.9 that

$$\mathcal{G} \stackrel{\text{def}}{=} \mathcal{O}_{S_{\mathbb{F}_p}} + \text{Image}\{(\Phi_{S_{\mathbb{F}_p}}^2)^*\left(\frac{1}{p} \cdot \mathcal{I}\right)_{\mathbb{F}_p}\} \subseteq (\mathcal{B}_{S'})_{\mathbb{F}_p}$$

is precisely the subsheaf of *polynomials of degree ≤ 1* arising from the torsor structure on $F'_{\mathbf{F}_p}$ (cf. the statement of Lemma 1.9). Note, moreover, that by Condition (2) of Definition 1.13 (i.e., the fact that the morphism $H \rightarrow (F')^{\text{st}}$ is *étale* – hence induces an *isomorphism* on PD-envelopes), it follows that

$$\mathcal{G} \subseteq (\mathcal{B}_{S'}^{\text{st}})_{\mathbf{F}_p} + (\mathcal{I}_{H_{B_{S'}}})_{\mathbf{F}_p} \subseteq (\mathcal{B}_{S'})_{\mathbf{F}_p}$$

(Indeed, this is a consequence of the fact that any section of $\text{Image}\{(\Phi_{S_{\mathbf{F}_p}}^2)^*(\frac{1}{p} \cdot \mathcal{I})_{\mathbf{F}_p}\}$ either already belongs to $(\mathcal{B}_{S'}^{\text{st}})_{\mathbf{F}_p}$, or *appears as the left-hand side* of an equation of the form

$$\text{linear term in weak variables} = \text{PD power series in strong variables}$$

which vanishes on the Hodge subspace; moreover, the existence of such an equation (which vanishes on the Hodge subspace) follows from the fact that the morphism $H \rightarrow (F')^{\text{st}}$ induces an isomorphism on PD-envelopes.)

Thus, putting the above observations together, we conclude that *the image of \mathcal{G} in $\mathcal{A}_{\mathbf{F}_p}$ (via the morphism $(\mathcal{B}_{S'})_{\mathbf{F}_p} \rightarrow \mathcal{A}_{\mathbf{F}_p}$) is contained in $(\mathcal{A}^{\text{st}})_{\mathbf{F}_p}$* . It thus follows that $\Phi_{\mathcal{I}}(\frac{1}{p} \cdot \mathcal{I})$ (where $\Phi_{\mathcal{I}}$ is as in the proof of Lemma 1.6) maps to zero modulo p in $\mathcal{E}_{\mathbf{F}_p} = (\mathcal{A}/\mathcal{A}^{\text{st}})_{\mathbf{F}_p}$. Since $\mathcal{I} \subseteq \mathcal{A}$ surjects onto \mathcal{E} , it thus follows that $\Phi_{\mathcal{E}}$ is divisible by p^2 , as desired. \circ

Let us write

$$\mathcal{E}^1 \subseteq \mathcal{E}$$

for the p -adic closure of the intersection of the images of the $(\frac{1}{p^2} \cdot \Phi_{\mathcal{E}})^N$ (for $N \geq 1$). Since $\Phi_{\mathcal{E}}$ is compatible with $\nabla_{\mathcal{E}}$, it follows that $\nabla_{\mathcal{E}}$ stabilizes \mathcal{E}^1 .

Lemma 2.8. *We have a natural horizontal isomorphism of vector bundles on S :*

$$\mathcal{E}^1 \cong \Omega_{S^{\text{log}}}^{\text{wk}}$$

(where the connection on $\Omega_{S^{\text{log}}}^{\text{wk}}$ is the one for which sections of $\Omega_{S^{\text{log}}}^{\text{wk,et}}$ are horizontal) that is compatible with the Frobenius actions given by $\frac{1}{p^2} \cdot \Phi_{\mathcal{E}}$ and $\Omega_{\Phi}^{\text{wk}}$.

Proof. The proof is entirely similar to that of Lemma 1.6. \circ

Next, let us define $\text{Geo}(\mathcal{A})$ by the fiber product:

$$\text{Geo}(\mathcal{A}) \stackrel{\text{def}}{=} \mathcal{A} \times_{\mathcal{E}} \mathcal{E}^1 \subseteq \mathcal{A}$$

Note that $\text{Geo}(\mathcal{A})$ is *preserved* by $\nabla_{\mathcal{A}}$ and $\Phi_{\mathcal{A}}$. Moreover, $\text{Geo}(\mathcal{A})$ fits into an *exact sequence*

$$0 \rightarrow \mathcal{A}^{\text{st}} \rightarrow \text{Geo}(\mathcal{A}) \rightarrow \Omega_{\mathcal{S}^{\log}}^{\text{wk}} \rightarrow 0$$

stabilized by $\nabla_{\mathcal{A}}$ and $\Phi_{\mathcal{A}}$. Thus, one may think of $\text{Geo}(\mathcal{A})$ as consisting of functions (i.e., sections of \mathcal{A}) of the form:

$$(\text{PD power series in strong variables}) + (\text{weak linear terms})$$

Next, we define

$$\text{Ind}(\mathcal{A}) \stackrel{\text{def}}{=} \text{Geo}(\mathcal{A}) \cdot \mathcal{A}^{\text{st}} \subseteq \mathcal{A}$$

Here, “Ind” stands for “induction” since we will use $\text{Ind}(\mathcal{A})$ in order to carry out the technique of *crystalline induction* (cf. [Mzk1], Chapter V) in §2.3. Thus, $\text{Ind}(\mathcal{A})$ is a *free module of rank* $1 + (d - c)$ *over* \mathcal{A}^{st} which fits into an exact sequence

$$0 \rightarrow \mathcal{A}^{\text{st}} \rightarrow \text{Ind}(\mathcal{A}) \rightarrow \Omega_{\mathcal{S}^{\log}}^{\text{wk}} \otimes_{\mathcal{O}_{\mathcal{S}}} \mathcal{A}^{\text{st}} \rightarrow 0$$

That is to say, $\text{Ind}(\mathcal{A})$ consists of:

polynomials of degree ≤ 1 in the weak variables with coefficients which are PD power series in the strong variables

One sees easily that $\text{Geo}(\mathcal{A})$ and $\text{Ind}(\mathcal{A})$ are preserved by $\nabla_{\mathcal{A}_T}$, hence inherit their own *connections* $\nabla_{\text{Geo}(\mathcal{A})}$ and $\nabla_{\text{Ind}(\mathcal{A})}$. Moreover, the Hodge filtration on \mathcal{A} induces *Hodge filtrations* on $\text{Geo}(\mathcal{A})$ and $\text{Ind}(\mathcal{A})$. Finally, by pulling back to T , we obtain $\text{Geo}(\mathcal{A}_T)$ and $\text{Ind}(\mathcal{A}_T)$.

Next, let us note that we have a morphism

$$\frac{1}{p} \cdot \Phi_{T_{\mathbb{F}_p}}^* (\mathcal{A}/\mathcal{A}^{\text{st}})_T = \frac{1}{p} \cdot \Phi_{T_{\mathbb{F}_p}}^* \mathcal{E}_T \rightarrow \mathcal{B}_T/\mathcal{B}_T^{\text{st}}$$

Let us denote the image of $\frac{1}{p} \cdot \Phi_{T_{\mathbb{F}_p}}^* (\mathcal{E}_T^1)$ via this morphism by

$$\text{Geo}(\mathcal{B}_T) + \mathcal{B}_T^{\text{st}} \subseteq \mathcal{B}_T/\mathcal{B}_T^{\text{st}}$$

and by

$$\mathrm{Geo}(\mathcal{B}_T) \subseteq \mathcal{B}_T$$

the inverse image of $\mathrm{Geo}(\mathcal{B}_T) + \mathcal{B}_T^{\mathrm{st}}$ in \mathcal{B}_T . Also, we shall write

$$\mathrm{Ind}(\mathcal{B}_T) \stackrel{\mathrm{def}}{=} \mathrm{Geo}(\mathcal{B}_T) \cdot \mathcal{B}_T^{\mathrm{st}} \subseteq \mathcal{B}_T$$

Thus, we have *exact sequences*:

$$0 \rightarrow \mathcal{B}_T^{\mathrm{st}} \rightarrow \mathrm{Geo}(\mathcal{B}_T) \rightarrow \Phi_{T_{\mathbb{F}_p}}^*(\Omega_{S^{\mathrm{log}}|T}^{\mathrm{wk}}) \rightarrow 0$$

and

$$0 \rightarrow \mathcal{B}_T^{\mathrm{st}} \rightarrow \mathrm{Ind}(\mathcal{B}_T) \rightarrow \Phi_{T_{\mathbb{F}_p}}^*(\Omega_{S^{\mathrm{log}}|T}^{\mathrm{wk}}) \otimes_{\mathcal{O}_T} \mathcal{B}_T^{\mathrm{st}} \rightarrow 0$$

In particular, $\mathrm{Ind}(\mathcal{B}_T)$ is a free $\mathcal{B}_T^{\mathrm{st}}$ -module of rank $1 + (d - c)$. Moreover, we get induced *connections* and *Hodge filtrations* on $\mathrm{Geo}(\mathcal{B}_T)$ and $\mathrm{Ind}(\mathcal{B}_T)$. Note that (since the morphism induced by $H \rightarrow (F')^{\mathrm{st}}$ on PD-envelopes is an isomorphism) the *Hodge filtration* on $\mathcal{B}_T^{\mathrm{st}}$ satisfies:

$$F^1(-) = 0; \quad F^0(-) = \mathcal{B}_T^{\mathrm{st}}$$

while the *Hodge filtration* on $\mathrm{Ind}(\mathcal{B}_T)$ satisfies:

$$F^2(-) = 0; \quad F^1(-) \cong \Phi_{T_{\mathbb{F}_p}}^*(\Omega_{S^{\mathrm{log}}|T}^{\mathrm{wk}}) \otimes_{\mathcal{O}_T} \mathcal{B}_T^{\mathrm{st}}; \quad F^0(-) = \mathrm{Ind}(\mathcal{B}_T)$$

(where the isomorphism arises from the above exact sequence).

Remark. Note that if we pull-back $\mathrm{Ind}(\mathcal{A})$ and $\mathrm{Ind}(\mathcal{B}_T)$ to $\widehat{S}^{\mathrm{st}}$ (so as to obtain $\mathrm{Ind}(\mathcal{A}_{\widehat{S}^{\mathrm{st}}})$ and $\mathrm{Ind}(\mathcal{B}_{\widehat{S}^{\mathrm{st}}})$), and then tensor with the augmentation morphisms $\mathcal{A}_{\widehat{S}^{\mathrm{st}}}^{\mathrm{st}} \rightarrow \mathcal{O}_{\widehat{S}^{\mathrm{st}}}$, $\mathcal{B}_{\widehat{S}^{\mathrm{st}}}^{\mathrm{st}} \rightarrow \mathcal{O}_{\widehat{S}^{\mathrm{st}}}$ (which are *horizontal* over $\widehat{S}^{\mathrm{st}}$), we obtain a pair of vector bundles

$$P_{\mathcal{A}} \stackrel{\mathrm{def}}{=} \mathrm{Ind}(\mathcal{A}_{\widehat{S}^{\mathrm{st}}}) \otimes_{\mathcal{A}_{\widehat{S}^{\mathrm{st}}}^{\mathrm{st}}} \mathcal{O}_{\widehat{S}^{\mathrm{st}}}; \quad P_{\mathcal{B}} \stackrel{\mathrm{def}}{=} \mathrm{Ind}(\mathcal{B}_{\widehat{S}^{\mathrm{st}}}) \otimes_{\mathcal{B}_{\widehat{S}^{\mathrm{st}}}^{\mathrm{st}}} \mathcal{O}_{\widehat{S}^{\mathrm{st}}}$$

of rank $1 + (d - c)$ on $\widehat{S}^{\mathrm{st}}$ equipped with *connections*, *Hodge filtrations*, and a *Frobenius action* (consisting of isomorphisms $\mathbf{F}^*(P_{\mathcal{A}}) \cong P_{\mathcal{B}}$; $\mathbf{F}^*(P_{\mathcal{B}}) \cong P_{\mathcal{A}}$). (Indeed, all of these structures are induced from the corresponding structures of the objects considered in Theorem 1.16.) Moreover,

sorting through the definitions reveals that the *inductive limit of the diagram* (equipped with its induced connection, Hodge filtration, and Frobenius action)

$$\begin{array}{ccc} \mathcal{O}_{\widehat{S}^{\text{st}}} \oplus \mathcal{O}_{\widehat{S}^{\text{st}}} & \longrightarrow & P_{\mathcal{A}} \oplus P_{\mathcal{B}} \\ \downarrow & & \\ \mathcal{O}_{\widehat{S}^{\text{st}}} & & \end{array}$$

(where the horizontal morphism is the natural morphism induced by the inclusions $\mathcal{A}^{\text{st}} \hookrightarrow \text{Ind}(\mathcal{A})$, $\mathcal{B}_T^{\text{st}} \hookrightarrow \text{Ind}(\mathcal{B}_T)$) *may be identified with the \mathcal{MF}^{∇} -object*

$$(P, \nabla_P, F^*(P), \Phi_P)_{(\widehat{S}^{\text{st}})^{\log}}^{\text{wk}}$$

considered at the end of §2.1.

Finally, before continuing, we note that the subsheaf

$$\text{Geo}(\mathcal{A}) \subseteq \mathcal{A}$$

constitutes (in the present context) the *geometrizing sub-object* discussed in Chapter VIII, §1 (which, unfortunately, in this case is not of finite rank). That is to say, the *zero loci* of the functions (i.e., sections of \mathcal{A}) that are contained in $\text{Geo}^{\text{st}}(\mathcal{A})$, $\text{Geo}(\mathcal{A}) \setminus \text{Geo}^{\text{st}}(\mathcal{A})$ will be the *geodesics* for this geometry. In particular, the geodesics will *either* be of the form

$$\text{linear expression in strong variables} = 0$$

or of the form

$$\begin{aligned} & \text{(nonzero) linear expression in weak variables} \\ & = \text{PD power series in strong variables} \end{aligned}$$

In other words, although the portion of the *geometry that relates to the strong variables is affine*, the portion of the *geometry that relates to the weak variables is non-affine*, i.e., it is affine, or linear, *only in the fibers over the quotient space defined by the strong variables*. We give an illustration of this sort of geometry in the Pictorial Appendix.

§2.3. Construction of the Galois Mantle: The Spiked Case

In this §, which constitutes the climax of our treatment of spiked Frobenius liftings, we would like to associate a *canonical Galois representation to the formal (infinite rank) \mathcal{MF}^{∇} -object of §1* (i.e., Theorem 1.16).

Unlike the strong portion of the mantle (cf. §1.2, 1.3), the total mantle (i.e., \mathcal{A} , which consists of both strong *and* weak parts) cannot be treated using the conventional finite rank techniques of [Falt1], §2. Instead, we must employ the technique of *crystalline induction* (cf. [Mzk1], Chapter V). (We remark, however, that the treatment that we give here does not presume a knowledge of [Mzk1], Chapter V.) The key idea is the following:

We regard the total mantle \mathcal{A} as an object *over* the strong mantle \mathcal{A}^{st} . Since \mathcal{A}^{st} already has an associated Galois object $\mathcal{A}^{\text{st}, \text{Gal}}$ (see below), if we work over the ring of p -adic periods, then we may *base-change* \mathcal{A} from \mathcal{A}^{st} to $\mathcal{A}^{\text{st}, \text{Gal}}$. After this base-change, we obtain an object over $\mathcal{A}^{\text{st}, \text{Gal}}$ (tensored with the ring of p -adic periods) which is *essentially of finite rank*, thus allowing us to associate to it a Galois representation via the theory of [Falt1], §2.

Finally, after we associate a Galois representation to \mathcal{A} , we discuss the relationship between this representation and the Galois representations obtained in §2.1.

Since our construction will be canonical, we may work étale locally on S . Thus, we may assume that S is *affine*. We may also assume that S^{\log} is *small* (in the sense of [Falt1], §2): that is, S^{\log} is *log étale over* (the p -adic completion of)

$$A[T_1, \dots, T_d]$$

(with the log structure given by the divisor $T_1 \cdots T_d$). We shall call these parameters T_1, \dots, T_d *small parameters*. Then we would like to consider the ring

$$B^+(S^{\log})$$

of [Falt1], §2. Recall that this ring is formed as follows:

- (1) First, we take the *normalization* \tilde{S}^{\log} of S^{\log} in the maximal covering of S_K^{\log} which is log étale in characteristic zero.
- (2) Since $\tilde{S}^{\log} \otimes \mathbf{F}_p$ is *epiperfect*, we can form $B(\tilde{S}^{\log} \otimes \mathbf{F}_p)$ as in Definition 1.13 of Chapter VI. This object $B(\tilde{S}^{\log} \otimes \mathbf{F}_p)$ is equipped with a structure of *PD-thickening*

$$\tilde{S}^{\log} \otimes \mathbf{F}_p \hookrightarrow B(\tilde{S}^{\log} \otimes \mathbf{F}_p)$$

of $\tilde{S}^{\log} \otimes \mathbf{F}_p$.

- (3) Then $B^+(S^{\log})$ is defined to be the ring whose Spf is the completion of $B(\tilde{S} \otimes \mathbf{F}_p)$ with respect to the topology defined by the divided powers of the ideal defining the embedding $\tilde{S}^{\log} \otimes \mathbf{F}_p \hookrightarrow B(\tilde{S}^{\log} \otimes \mathbf{F}_p)$.

We remark that in fact, it is sufficient to work with the smaller ring whose Spf is just $B(\tilde{S} \otimes \mathbf{F}_p)$ itself (i.e., *without passing to the completion*) – cf. Chapter VI, Example 1.8. Since, however, [Falt1] uses the completed ring, and we wish to quote the theory of [Falt1] in the following discussion, we shall also use the completed ring here, so that the results of [Falt1] are directly (as opposed to just “almost”) quotable.

Note, in particular, that

- (1) $B^+(S^{\log})$ is obtained as the inverse limit of a *projective system of PD-thickenings* of the $\mathcal{O}_{S_{\mathbf{F}_p}}$ -algebra $\mathcal{O}_{\tilde{S}_{\mathbf{F}_p}}$.
- (2) $B^+(S^{\log})$ has a PD-ideal

$$I^+ \subseteq B^+(S^{\log})$$

which is Galois-invariant and such that $B^+(S^{\log})/I^+ \cong \hat{\mathcal{O}}_{\tilde{S}}$ (i.e., the p -adic completion of $\mathcal{O}_{\tilde{S}}$).

Moreover, $B^+(S^{\log})$ comes equipped with a natural *Frobenius action* (base-change by which we shall denote by means of a superscripted “ F ”), as well as a continuous *Galois action*, i.e., an action of

$$\Pi_{S^{\log}} \stackrel{\text{def}}{=} \pi_1(S_K^{\log}, s)$$

which commutes with the Frobenius action. The Frobenius invariants of $B^+(S^{\log})$ are given by $\mathbf{Z}_p \subseteq B^+(S^{\log})$. There is a *Galois equivariant injection*

$$\beta : \mathbf{Z}_p(1) \hookrightarrow B^+(S^{\log})$$

The Frobenius action on β takes β to $p \cdot \beta$. This completes our review of $B^+(-)$ (cf. [Falt1], §2, for more details).

Now let us return to our discussion of *spiked Frobenius liftings*. First, recall the exact sequence of $\Pi_{S^{\log}}$ -modules of the discussion preceding

Definition 1.4 (derived by considering the *strong* portion of the uniformization). By pushing forward this exact sequence by $\mathbf{Z}_p \hookrightarrow W(\mathbf{F}_{p^2})$, we get an exact sequence

$$0 \rightarrow W(\mathbf{F}_{p^2}) \rightarrow (G^{\text{st}})_{\text{Gal}}^{\mathbf{F}_{p^2}} \rightarrow (\Omega_S^{\text{st,et}}) \otimes_{W(\mathbf{F}_{p^2})} (T_p(G_2)) \rightarrow 0$$

By considering $W(\mathbf{F}_{p^2})$ -linear splittings of this exact sequence, we obtain a *torsor* with Galois action. This torsor is the Spf of a p -adically complete ring with Galois action, which we shall refer to as the *affinization* associated to the above exact sequence (cf. Definition 1.7). Thus, the affinization

$$\text{Aff}((G^{\text{st}})_{\text{Gal}}^{\mathbf{F}_{p^2}})$$

is a topological ring with a continuous $\Pi_{S^{\log}}$ -action which is non-canonically (i.e., non-Galois equivariantly) isomorphic to (the p -adic completion of) a *polynomial algebra* of dimension c over $W(\mathbf{F}_{p^2})$. Moreover, the submodule of this algebra consisting of polynomials of degree ≤ 1 may be identified with

$$(G^{\text{st}})_{\text{Gal}}^{\mathbf{F}_{p^2}} \subseteq \text{Aff}((G^{\text{st}})_{\text{Gal}}^{\mathbf{F}_{p^2}})$$

If we *restrict* this representation to

$$\Pi_{T^{\log}} \subseteq \Pi_{S^{\log}}$$

then *modulo* p , there is a unique $\Pi_{T^{\log}}$ -invariant augmentation

$$\alpha_{\text{Gal}} : \text{Aff}((G^{\text{st}})_{\text{Gal}}^{\mathbf{F}_{p^2}}) \rightarrow \mathbf{F}_{p^2}$$

Let us denote the p -adic completion of the PD-envelope of $\text{Aff}((G^{\text{st}})_{\text{Gal}}^{\mathbf{F}_{p^2}})$ at this augmentation by

$$\mathcal{A}^{\text{st,Gal}}$$

Thus, $\mathcal{A}^{\text{st,Gal}}$ is a ring with $\Pi_{T^{\log}}$ -action.

Next, let us note that

$$\text{Geo}^{\text{st}}(\mathcal{A}) \oplus \text{Geo}^{\text{st}}(\mathcal{F}')$$

has the natural structure of an object of $\mathcal{MF}^{\nabla}(S^{\log})$ whose (covariantly! – since the original exact sequence preceding Definition 1.4

is obtained by taking the *dual* of the corresponding sequence of *contravariantly* associated p -divisible groups) *associated Galois representation* is precisely $(G^{\text{st}})_{\text{Gal}}^{\mathbf{F}_{p^2}}$. In particular, $\text{Geo}^{\text{st}}(\mathcal{A})$ and $\text{Geo}^{\text{st}}(\mathcal{F}')$ define *crystals* which we may evaluate on the *PD-thickening* $B^+(T^{\log})$ to form finite, free $B^+(T^{\log})$ -modules $\text{Geo}^{\text{st}}(\mathcal{A})(B^+(T^{\log}))$, $\text{Geo}^{\text{st}}(\mathcal{F}')(B^+(T^{\log}))$. Moreover, it follows from the way associated Galois representations are constructed in [Falt1], §2, that we have natural $\Pi_{T^{\log}}$ - and *Frobenius-equivariant, Hodge filtration-preserving, $W(\mathbf{F}_{p^2})$ -linear morphisms*

$$\{(G^{\text{st}})_{\text{Gal}}^{\mathbf{F}_{p^2}}\}^{\vee} \rightarrow \text{Geo}^{\text{st}}(\mathcal{A})(B^+(T^{\log}))^{\vee}$$

$$\{((G^{\text{st}})_{\text{Gal}}^{\mathbf{F}_{p^2}})^F\}^{\vee} \rightarrow \text{Geo}^{\text{st}}(\mathcal{F}')(B^+(T^{\log}))^{\vee}$$

(where the “ \vee ’s” on the left (respectively, right) are the duals over $W(\mathbf{F}_{p^2})$ (respectively, $B^+(T^{\log}) = B^+(S^{\log})$)). If we *tensor* these two morphisms over $W(\mathbf{F}_{p^2})$ with $B^+(T^{\log})$ on the *left* so as to obtain $B^+(T^{\log})$ -linear morphisms, and then form the *duals* (over $B^+(T^{\log})$) of these morphisms, we obtain natural $\Pi_{T^{\log}}$ - and *Frobenius-equivariant, Hodge filtration-preserving, $B^+(T^{\log})$ -linear morphisms*

$$(\text{Geo}^{\text{st}}(\mathcal{A}))(B^+(T^{\log})) \rightarrow B^+(T^{\log}) \otimes_{W(\mathbf{F}_{p^2})} (G^{\text{st}})_{\text{Gal}}^{\mathbf{F}_{p^2}} \hookrightarrow B^+(T^{\log}) \otimes_{W(\mathbf{F}_{p^2})} \mathcal{A}^{\text{st,Gal}}$$

$$\begin{aligned} (\text{Geo}^{\text{st}}(\mathcal{F}'))(B^+(T^{\log})) &\rightarrow B^+(T^{\log}) \otimes_{W(\mathbf{F}_{p^2})} ((G^{\text{st}})_{\text{Gal}}^{\mathbf{F}_{p^2}})^F \\ &\hookrightarrow B^+(T^{\log}) \otimes_{W(\mathbf{F}_{p^2})} (\mathcal{A}^{\text{st,Gal}})^F \end{aligned}$$

(where we understand the tensor products on the right to be *topological*, i.e., with respect to the p -adic topology). Moreover, by the universal properties of affinization and PD-envelopes (note that $\mathcal{A}_T^{\text{st}}$ (respectively, $\mathcal{B}_T^{\text{st}}$) may be naturally identified with the p -adic completion of a certain PD-envelope of the affinization of $\text{Geo}^{\text{st}}(\mathcal{A})_T$ (respectively, $\text{Geo}^{\text{st}}(\mathcal{F}')_T$) – cf. §1.3, 2.2), these morphisms induce morphisms

$$(\mathcal{A}_T^{\text{st}})(B^+(T^{\log})) \rightarrow \mathcal{A}_{B^+}^{\text{st,Gal}}$$

and

$$(\mathcal{B}_T^{\text{st}})(B^+(T^{\log})) \rightarrow (\mathcal{A}_{B^+}^{\text{st,Gal}})^F$$

where

$$\mathcal{A}_{B^+}^{\text{st,Gal}} \stackrel{\text{def}}{=} B^+(T^{\log}) \otimes_{W(\mathbf{F}_{p^2})} \mathcal{A}^{\text{st,Gal}}$$

(and the tensor product is *topological*, i.e., with respect to the p -adic topology). These morphisms are, of course, *compatible with $\Pi_{T^{\log}}$, the Frobenius action, and the Hodge filtration*. (Throughout, we regard the Galois representation $\mathcal{A}^{\text{st}, \text{Gal}}$ (as well as all other Galois representations) as equipped with the *trivial* Frobenius action and Hodge filtration.) It is relative to these two morphisms that we will perform the crucial *base-change* advertized at the beginning of this subsection.

Now let us recall that the crystal $(\text{Ind}(\mathcal{A}_T), \nabla_{\text{Ind}(\mathcal{A}_T)})$ on $\text{Crys}(T_{\mathbf{F}_p}^{\log})$ (cf. §2.2) is a *locally free $\mathcal{A}_T^{\text{st}}$ -module* of rank $1 + (d - c)$. Thus, if we *base-change* the $(\mathcal{A}_T^{\text{st}})(B^+(T^{\log}))$ -module

$$(\text{Ind}(\mathcal{A}_T)(B^+(T^{\log})))^{\vee}$$

(where the “ \vee ” denotes the $(\mathcal{A}_T^{\text{st}})(B^+(T^{\log}))$ -dual) by the above morphism

$$(\mathcal{A}_T^{\text{st}})(B^+(T^{\log})) \rightarrow \mathcal{A}_{B^+}^{\text{st}, \text{Gal}}$$

we get a *free $\mathcal{A}_{B^+}^{\text{st}, \text{Gal}}$ -module*

$$\text{Ind}^{\text{Gal}}(\mathcal{A})$$

of rank $1 + (d - c)$. Similarly, we may construct from $(\text{Ind}(\mathcal{B}_T)(B^+(T^{\log})))^{\vee}$ (cf. §2.2) a *free $(\mathcal{A}_{B^+}^{\text{st}, \text{Gal}})^F$ -module*

$$\text{Ind}^{\text{Gal}}(\mathcal{B})$$

of rank $1 + (d - c)$. Moreover, $\text{Ind}^{\text{Gal}}(\mathcal{A})$ and $\text{Ind}^{\text{Gal}}(\mathcal{B})$ come equipped with $\Pi_{T^{\log}}$ - and *Frobenius actions and Hodge filtrations $F(-)$* , and fit into *exact sequences*

$$0 \rightarrow \mathcal{A}_{B^+}^{\text{st}, \text{Gal}} \otimes_{W(\mathbf{F}_{p^2})} \Theta_{S^{\log}}^{\text{wk}, \text{et}}[-1] \rightarrow \text{Ind}^{\text{Gal}}(\mathcal{A}) \rightarrow \mathcal{A}_{B^+}^{\text{st}, \text{Gal}} \rightarrow 0$$

$$0 \rightarrow (\mathcal{A}_{B^+}^{\text{st}, \text{Gal}} \otimes_{W(\mathbf{F}_{p^2})} \Theta_{S^{\log}}^{\text{wk}, \text{et}})^F[-1] \rightarrow \text{Ind}^{\text{Gal}}(\mathcal{B}) \rightarrow (\mathcal{A}_{B^+}^{\text{st}, \text{Gal}})^F \rightarrow 0$$

that respect all of these structures. Here, the “ $[-1]$ ” denotes a -1 *shift in the Hodge filtration* (cf. the discussion of the Hodge filtration on $\text{Ind}(\mathcal{A}_T)$, $\text{Ind}(\mathcal{B}_T)$ in §2.2).

Next, recall that we have an *augmentation*

$$B^+(T^{\log})_{\mathbf{F}_p} \rightarrow (B^+(T^{\log})/I^+)_{\mathbf{F}_p} = \mathcal{O}_{\tilde{S}_{\mathbf{F}_p}}$$

whose kernel is annihilated by Frobenius (since I^+ has divided powers). Note, moreover, that

$$F^1(\mathcal{A}_{B^+}^{\text{st}, \text{Gal}} \otimes_{B^+(T^{\log})} \mathcal{O}_{\tilde{S}_{\mathbf{F}_p}}) = 0$$

(since the Hodge filtration of $\mathcal{A}_{B^+}^{\text{st}, \text{Gal}}$ is, by definition, *trivial*, while the Hodge filtration on $B^+(T^{\log})$ is that defined by the divided powers of I^+). Thus, it follows from the above exact sequences that:

$$F^1(\text{Ind}^{\text{Gal}}(\mathcal{A}) \otimes_{B^+(T^{\log})} \mathcal{O}_{\tilde{S}_{\mathbf{F}_p}}) = F^1(\text{Ind}^{\text{Gal}}(\mathcal{B}) \otimes_{B^+(T^{\log})} \mathcal{O}_{\tilde{S}_{\mathbf{F}_p}}) = 0$$

while

$$F^0(\text{Ind}^{\text{Gal}}(\mathcal{A}) \otimes_{B^+(T^{\log})} \mathcal{O}_{\tilde{S}_{\mathbf{F}_p}}) \quad \text{and} \quad F^0(\text{Ind}^{\text{Gal}}(\mathcal{B}) \otimes_{B^+(T^{\log})} \mathcal{O}_{\tilde{S}_{\mathbf{F}_p}})$$

are *free of rank 1* over $\mathcal{O}_{\tilde{S}_{\mathbf{F}_p}} \otimes_{W(\mathbf{F}_{p^2})} \mathcal{A}^{\text{st}, \text{Gal}}$ and $\mathcal{O}_{\tilde{S}_{\mathbf{F}_p}} \otimes_{W(\mathbf{F}_{p^2})} (\mathcal{A}^{\text{st}, \text{Gal}})^F$, respectively.

It thus follows that if we *base-change* $\text{Ind}^{\text{Gal}}(\mathcal{A})$ by the *Frobenius* on $B^+(T^{\log})$ to obtain $(\text{Ind}^{\text{Gal}}(\mathcal{A}))^F$ (an object whose reduction modulo p depends, by the above discussion, *only on* $\text{Ind}^{\text{Gal}}(\mathcal{A}) \otimes_{B^+(T^{\log})} \mathcal{O}_{\tilde{S}_{\mathbf{F}_p}}$), and then take the *submodule of sections whose reductions modulo p lie in $F^0(-)$* , we obtain a *free* $(\mathcal{A}_{B^+}^{\text{st}, \text{Gal}})^F$ -*module* (i.e., free by the above discussion, despite the “ \mathbf{F}^* ”)

$$\mathbf{F}^*(\text{Ind}^{\text{Gal}}(\mathcal{A}))$$

of rank $1 + (d - c)$. Similarly, we may construct the free $(\mathcal{A}_{B^+}^{\text{st}, \text{Gal}})$ -module

$$\mathbf{F}^*(\text{Ind}^{\text{Gal}}(\mathcal{B}))$$

(of rank $1 + (d - c)$). Moreover, the Frobenius action Ξ_{Φ} of Theorem 1.16 induces isomorphisms

$$\mathbf{F}^*(\text{Ind}^{\text{Gal}}(\mathcal{A})) \cong \text{Ind}^{\text{Gal}}(\mathcal{B}), \quad \mathbf{F}^*(\text{Ind}^{\text{Gal}}(\mathcal{B})) \cong \text{Ind}^{\text{Gal}}(\mathcal{A})$$

Thus, we obtain a *Frobenius action* on

$$\mathcal{M} \stackrel{\text{def}}{=} \text{Ind}^{\text{Gal}}(\mathcal{A}) \oplus \text{Ind}^{\text{Gal}}(\mathcal{B})$$

which respects the exact sequences discussed above.

Let us write

$$\hat{\mathcal{A}}^{\text{st}, \text{Gal}}$$

for the completion of $\mathcal{A}^{\text{st}, \text{Gal}}$ with respect to the topology defined by the divided powers of the kernel of the augmentation $\mathcal{A}^{\text{st}, \text{Gal}} \rightarrow \mathbf{F}_{p^2}$. Thus, $\hat{\mathcal{A}}^{\text{st}, \text{Gal}}$ is a profinite ring with continuous $\Pi_{T^{\log}}$ -action. Note, moreover, that the natural map $\mathcal{A}^{\text{st}, \text{Gal}} \rightarrow \hat{\mathcal{A}}^{\text{st}, \text{Gal}}$ induces inclusions

$$\mathcal{A}^{\text{st}, \text{Gal}} \otimes \mathbf{Z}/p^{n+1}\mathbf{Z} \hookrightarrow \hat{\mathcal{A}}^{\text{st}, \text{Gal}} \otimes \mathbf{Z}/p^{n+1}\mathbf{Z}$$

for all $n \geq 1$. Now we have the following important pair of technical lemmas:

Lemma 2.9. *Let \mathcal{R} be a profinite (commutative) ring, equipped with a continuous inclusion*

$$\hat{\mathcal{A}}^{\text{st}, \text{Gal}} \hookrightarrow \mathcal{R}$$

Suppose, moreover, that for every member $\mathcal{R} \rightarrow \mathcal{R}'$ of some cofinal (in an inverse system whose limit is \mathcal{R}) collection of finite $\mathcal{A}^{\text{st}, \text{Gal}}$ -algebra quotients $\{\mathcal{R} \rightarrow \mathcal{R}'\}$, the submodule

$$F^0(\mathcal{M} \otimes_{\mathcal{A}^{\text{st}, \text{Gal}}} \mathcal{R}')^{F=1} \subseteq F^0(\mathcal{M} \otimes_{\mathcal{A}^{\text{st}, \text{Gal}}} \mathcal{R}')$$

(of elements fixed by Frobenius) forms a free \mathcal{R}' -module of rank $1 + (d - c)$. Then it follows that the submodule

$$G_{\Phi}^{\text{wk}} \stackrel{\text{def}}{=} F^0(\mathcal{M})^{F=1} \subseteq F^0(\mathcal{M})$$

forms a free $\mathcal{A}^{\text{st}, \text{Gal}}$ -module of rank $1 + (d - c)$. Moreover, for \mathcal{C} equal to any finite quotient of $\hat{\mathcal{A}}^{\text{st}, \text{Gal}}$, the module $F^0(\mathcal{M} \otimes_{\mathcal{A}^{\text{st}, \text{Gal}}} \mathcal{C})^{F=1}$ is free over \mathcal{C} of rank $1 + (d - c)$ and generated (over \mathcal{C}) by the image of $F^0(\mathcal{M})^{F=1}$.

Proof. First, observe that to show that $F^0(\mathcal{M})^{F=1}$ is free over $\mathcal{A}^{\text{st}, \text{Gal}}$ of rank $1 + (d - c)$, it suffices to show that, for each $n \geq 1$,

$$F^0(\mathcal{M} \otimes \mathbf{Z}/p^{n+1}\mathbf{Z})^{F=1}$$

is a free $(\mathcal{A}^{\text{st,Gal}} \otimes \mathbf{Z}/p^{n+1}\mathbf{Z})$ -module of rank $1 + (d - c)$. But note that since the topology on $\mathcal{A}^{\text{st,Gal}} \otimes \mathbf{Z}/p^{n+1}\mathbf{Z}$ is discrete, there exists a sub- $W(\mathbf{F}_{p^2})$ -algebra

$$\mathcal{K} \subseteq \mathcal{A}_{\mathbf{Z}/p^{n+1}\mathbf{Z}}^{\text{st,Gal}} \stackrel{\text{def}}{=} \mathcal{A}^{\text{st,Gal}} \otimes \mathbf{Z}/p^{n+1}\mathbf{Z}$$

which is *finitely generated over* $W(\mathbf{F}_{p^2})$, and is such that the $(\mathcal{A}_{B^+}^{\text{st,Gal}})_{\mathbf{Z}/p^{n+1}\mathbf{Z}}$ -module $\mathcal{M} \otimes \mathbf{Z}/p^{n+1}\mathbf{Z}$ is obtained as

$$\mathcal{M} \otimes \mathbf{Z}/p^{n+1}\mathbf{Z} = \mathcal{N} \otimes_{\mathcal{K}} \mathcal{A}_{\mathbf{Z}/p^{n+1}\mathbf{Z}}^{\text{st,Gal}} = \mathcal{N} \otimes_{\mathcal{K}_{B^+}} (\mathcal{A}_{B^+}^{\text{st,Gal}})_{\mathbf{Z}/p^{n+1}\mathbf{Z}}$$

for some \mathcal{K}_{B^+} -module \mathcal{N} (where the subscripted B^+ denotes the (p -adic) topological tensor product over $W(\mathbf{F}_{p^2})$ with $B^+(T^{\log})$). We may even assume that \mathcal{N} is a free \mathcal{K}_{B^+} -module of rank $1 + (d - c)$ (since $\mathcal{M} \otimes \mathbf{Z}/p^{n+1}\mathbf{Z}$ is a free $(\mathcal{A}_{B^+}^{\text{st,Gal}})_{\mathbf{Z}/p^{n+1}\mathbf{Z}}$ -module of the same rank), and that the *Hodge filtration* and *Frobenius action* on $\mathcal{M} \otimes \mathbf{Z}/p^{n+1}\mathbf{Z}$ arise from a Hodge filtration and Frobenius action on \mathcal{N} . In particular, it makes sense to consider the \mathcal{K} -submodule

$$F^0(\mathcal{N})^{F=1} \subseteq F^0(\mathcal{N}) \subseteq \mathcal{N}$$

Note that if we can show that $F^0(\mathcal{N})^{F=1}$ is free over \mathcal{K} of rank $1 + (d - c)$, then it will follow that $F^0(\mathcal{M})^{F=1}$ is free over $\mathcal{A}^{\text{st,Gal}}$ of rank $1 + (d - c)$.

Next, let us observe that it follows from the definition of $\mathcal{A}_{\mathbf{Z}/p^{n+1}\mathbf{Z}}^{\text{st,Gal}}$ that the kernel of the augmentation $\mathcal{A}_{\mathbf{Z}/p^{n+1}\mathbf{Z}}^{\text{st,Gal}} \rightarrow \mathbf{F}_{p^2}$ consists of *nilpotent elements*. Thus, it follows that the ring \mathcal{K} also admits a surjection

$$\mathcal{K} \rightarrow \mathbf{F}_{p^2}$$

whose kernel consists of nilpotent elements. Since \mathcal{K} is finitely generated over $W(\mathbf{F}_{p^2})$, this implies that \mathcal{K} is a *local artinian ring* whose underlying set is *finite*. On the other hand, we have

$$\mathcal{K} \subseteq \mathcal{A}_{\mathbf{Z}/p^{n+1}\mathbf{Z}}^{\text{st,Gal}} \subseteq \widehat{\mathcal{A}}_{\mathbf{Z}/p^{n+1}\mathbf{Z}}^{\text{st,Gal}} \subseteq \mathcal{R}/(p^{n+1} \cdot \widehat{\mathcal{A}}^{\text{st,Gal}})$$

Moreover, since the inclusion $\widehat{\mathcal{A}}^{\text{st,Gal}} \hookrightarrow \mathcal{R}$ is *continuous* and $\widehat{\mathcal{A}}^{\text{st,Gal}}$ is profinite, hence *compact*, it follows that $p^{n+1} \cdot \widehat{\mathcal{A}}^{\text{st,Gal}}$ is *closed* in \mathcal{R} . Since \mathcal{R} is profinite, it thus follows that (the $\widehat{\mathcal{A}}^{\text{st,Gal}}$ -module) $\mathcal{R}/(p^{n+1} \cdot \widehat{\mathcal{A}}^{\text{st,Gal}})$ is also *profinite*, hence that there exists a *quotient*

$$\mathcal{R}/(p^{n+1} \cdot \hat{\mathcal{A}}^{\text{st}, \text{Gal}}) \rightarrow \mathcal{Q}$$

of $\hat{\mathcal{A}}^{\text{st}, \text{Gal}}$ -modules such that the natural morphism $\mathcal{K} \rightarrow \mathcal{Q}$ is an *inclusion*. (Here we use the fact that the underlying set of \mathcal{K} is *finite*.) Note that (by the hypotheses of Lemma 2.9) there exists a *quotient* $\hat{\mathcal{A}}^{\text{st}, \text{Gal}}$ -algebra

$$\mathcal{R} \rightarrow \mathcal{R}'$$

which is a *member of the cofinal collection* of the statement of Lemma 2.9, and which is such that the quotient $\mathcal{R} \rightarrow \mathcal{Q}$ *factors through* \mathcal{R}' .

Now let us recall our assumption that

$$F^0(\mathcal{M} \otimes_{\mathcal{A}^{\text{st}, \text{Gal}}} \mathcal{R}')^{F=1} \subseteq F^0(\mathcal{M} \otimes_{\mathcal{A}^{\text{st}, \text{Gal}}} \mathcal{R}')$$

is a free \mathcal{R}' -module of rank $1 + (d - c)$. On the other hand, recall (cf. [Falt1], §2, g)) that for any $\mathcal{A}^{\text{st}, \text{Gal}}$ -module \mathcal{E} whose *cardinality* (as a set) $|\mathcal{E}|$ is $< \infty$, we have

$$|F^0(\mathcal{M} \otimes_{\mathcal{A}^{\text{st}, \text{Gal}}} \mathcal{E})^{F=1}| \leq |\mathcal{E}|^{1+(d-c)}$$

(Indeed, this follows by “*dévissage*” – which allows one to reduce to the case $\mathcal{E} = \mathbf{F}_{p^2}$; but in this case, $\mathcal{M} \otimes_{\mathcal{A}^{\text{st}, \text{Gal}}} \mathbf{F}_{p^2}$ is essentially just the (reduction modulo p) of the “canonical weak uniformizing \mathcal{MF}^∇ -object” of the discussion proceeding Definition 2.5 (cf. also the Remark of §2.2), hence we may conclude (not just “ \leq ” but strict equality) by applying [Falt1], §2, Theorem 2.4, directly.) Thus, we conclude that if \mathcal{E} is the kernel of the surjection $\mathcal{R}' \rightarrow \mathcal{Q}$, then we have an exact sequence of finite $W(\mathbf{F}_{p^2})$ -modules

$$0 \rightarrow F^0(\mathcal{M} \otimes_{\mathcal{A}^{\text{st}, \text{Gal}}} \mathcal{E})^{F=1} \rightarrow F^0(\mathcal{M} \otimes_{\mathcal{A}^{\text{st}, \text{Gal}}} \mathcal{R}')^{F=1} \rightarrow F^0(\mathcal{M} \otimes_{\mathcal{A}^{\text{st}, \text{Gal}}} \mathcal{Q})^{F=1} \rightarrow 0$$

i.e., in particular, the *cardinality* of $F^0(\mathcal{M} \otimes_{\mathcal{A}^{\text{st}, \text{Gal}}} \mathcal{Q})^{F=1}$ is equal to $|\mathcal{Q}|^{1+(d-c)}$.

Similarly, I *claim* that the sequence

$$0 \rightarrow F^0(\mathcal{N})^{F=1} \rightarrow F^0(\mathcal{N} \otimes_{\mathcal{K}} \mathcal{Q})^{F=1} \rightarrow F^0(\mathcal{N} \otimes_{\mathcal{K}} (\mathcal{Q}/\mathcal{K}))^{F=1} \rightarrow 0$$

is *exact*. Indeed, the middle term is the same as $F^0(\mathcal{M} \otimes_{\mathcal{A}^{\text{st}, \text{Gal}}} \mathcal{Q})^{F=1}$, hence has cardinality equal to $|\mathcal{Q}|^{1+(d-c)}$. Moreover, since \mathcal{K} is local artinian, one sees by the same “*dévissage*” argument as in the preceding paragraph that the *first and last terms* have cardinality no more than $|\mathcal{K}|^{1+(d-c)}$

and $|\mathcal{Q}/\mathcal{K}|^{1+(d-c)}$, respectively. This is sufficient to prove the claim that the sequence is exact. In particular, we conclude that $F^0(\mathcal{N})^{F=1}$ has the same cardinality as the free \mathcal{K} -module of rank $1 + (d - c)$. But this is only possible if $F^0(\mathcal{N})^{F=1}$ is free over \mathcal{K} of rank $1 + (d - c)$.

This completes the proof that G_{Φ}^{wk} is free over $\mathcal{A}^{\text{st}, \text{Gal}}$ of rank $1 + (d - c)$. Moreover, (relative to the last sentence in the statement of Lemma 2.9) by choosing n and \mathcal{K} so that the morphism $\mathcal{K} \rightarrow \mathcal{C}$ is surjective, we see that a “dévissage” argument as in the preceding paragraph implies the truth of the last sentence in the statement of Lemma 2.9. Thus, the proof of Lemma 2.9 is complete. \circ

Lemma 2.10. *There exists a ring \mathcal{R} satisfying the conditions of Lemma 2.9.*

Proof. The idea here is to embed the ring $\hat{\mathcal{A}}^{\text{st}, \text{Gal}}$ – which one may think of as a certain space of p -adic analytic functions on $\text{Spf}(\text{Aff}((G^{\text{st}})_{\text{Gal}}^{\mathbf{F}_{p^2}}))$ – in the space of set-theoretic functions on (a certain subset of) the set of points of $\text{Spf}(\text{Aff}((G^{\text{st}})_{\text{Gal}}^{\mathbf{F}_{p^2}}))$. Moreover, the restriction of \mathcal{M} to each of these points will be sufficiently close to a finite rank \mathcal{MF}^{∇} -object for us to apply the theory of [Falt1], §2.

We make this idea precise as follows: Recall the augmentation

$$\alpha_{\text{Gal}} : \text{Aff}((G^{\text{st}})_{\text{Gal}}^{\mathbf{F}_{p^2}}) \rightarrow \mathbf{F}_{p^2}$$

used to define $\mathcal{A}^{\text{st}, \text{Gal}}$ as (the p -adic completion of) a certain PD-envelope. Let

$$E \subseteq \text{Spf}(\text{Aff}((G^{\text{st}})_{\text{Gal}}^{\mathbf{F}_{p^2}}))(W(\mathbf{F}_{p^2}))$$

be the subset of all $W(\mathbf{F}_{p^2})$ -valued points whose reduction modulo p is equal to α_{Gal} . Thus, E is naturally a torsor over the free $W(\mathbf{F}_{p^2})$ -module (of rank c)

$$p \cdot \{(\Omega_S^{\text{st}, \text{et}} \otimes_{W(\mathbf{F}_{p^2})} (T_p(G_2)))\}^{\vee}$$

(where the “ \vee ” denotes the $W(\mathbf{F}_{p^2})$ -linear dual). Moreover, E is equipped with a natural $\Pi_{T^{\log}}$ -action, together with $\Pi_{T^{\log}}$ -equivariant natural inclusions:

$$E \subseteq \text{Spf}(\mathcal{A}^{\text{st}, \text{Gal}})(W(\mathbf{F}_{p^2})) \subseteq \text{Spf}(\hat{\mathcal{A}}^{\text{st}, \text{Gal}})(W(\mathbf{F}_{p^2}))$$

In particular, it follows that if we let

$$\mathcal{R} \stackrel{\text{def}}{=} \prod_E W(\mathbf{F}_{p^2})$$

(i.e., the product of copies of $W(\mathbf{F}_{p^2})$ – one for each element of E – equipped with the product of the topologies of the $W(\mathbf{F}_{p^2})$) be the *topological ring of set-theoretic functions* $E \rightarrow W(\mathbf{F}_{p^2})$, then, by *evaluation*, we get a natural *continuous ring homomorphism*

$$\rho : \widehat{\mathcal{A}}^{\text{st}, \text{Gal}} \rightarrow \mathcal{R}$$

It is easy to see that, moreover, that ρ is *injective*. Indeed, this amounts to the assertion that if an element ϕ of the completed PD-envelope at the origin (i.e., the point $X_1 = \dots = X_c = 0$) of $W(\mathbf{F}_{p^2})[X_1, \dots, X_c]$ satisfies $\phi(x_1, \dots, x_c) = 0$ for all $(x_1, \dots, x_c) \in p \cdot W(\mathbf{F}_{p^2})^c$, then $\phi \equiv 0$. But this is an easy exercise.

To complete the proof of Lemma 2.10, it suffices to show that for each point $e \in E$, the restricted object

$$\mathcal{M}_e \stackrel{\text{def}}{=} \mathcal{M} \otimes_{\mathcal{A}^{\text{st}, \text{Gal}}, e} W(\mathbf{F}_{p^2})$$

(where we think of e as a morphism $\mathcal{A}^{\text{st}, \text{Gal}} \rightarrow W(\mathbf{F}_{p^2})$) satisfies (for all $n \geq 1$):

(*) $F^0(\mathcal{M}_e \otimes \mathbf{Z}/p^{n+1}\mathbf{Z})^{F=1}$ is a free $W(\mathbf{F}_{p^2})_{\mathbf{Z}/p^{n+1}\mathbf{Z}}$ -module of rank $1 + (d - c)$.

First, let us observe that since the action of $\Pi_{T^{\log}}$ preserves the Hodge filtration and Frobenius action, it follows that for any $\gamma \in \Pi_{T^{\log}}$, the truth of (*) for e implies the truth of (*) for $\gamma \cdot e \in E$. Thus, it follows from the torsor structure of E and the fact that $\Pi_{T^{\log}}$ acts *transitively on* E (cf. Theorem 1.5 plus the fact that the covering $(\Phi^{\log})_K^n : S_K^{\log} \rightarrow S_K^{\log}$ is *connected*) that *once one shows (*) for some particular $e_0 \in E$, it follows that (*) holds for arbitrary $e \in E$.*

On the other hand, sorting through the definitions shows that the infinite log étale covering $((S^{\text{st}})^{\log})_K \rightarrow T_K^{\log}$ defines a *tautological* $e_0 \in E$ over (i.e., after restricting everything to) $(S^{\text{st}})^{\log}$. Moreover, the resulting object

$$\mathcal{M}_{e_0}$$

is essentially the same \mathcal{MF}^{∇} -object as the object appearing in the discussion preceding Definition 2.5 (cf. also the Remark of §2.2). (Here, meaning of “essentially” is the following: In the notation of the Remark of §2.2, \mathcal{M}_{e_0} corresponds to $P_{\mathcal{A}} \oplus P_{\mathcal{B}}$, as opposed to the

inductive limit of the diagram appearing in this Remark.) Thus, the truth of (*) for this \mathcal{M}_{e_0} follows directly from [Falt1], §2, Theorem 2.4. \bigcirc

Thus, we obtain the following result, which is the main result of this Chapter:

Theorem 2.11. *Let $\Phi^{\log} : S^{\log} \rightarrow S^{\log}$ be a very ordinary spiked Frobenius lifting (cf. Definition 1.1) of colevel c , equipped with a Hodge subspace H (cf. Definition 1.13). Then associated to Φ^{\log} , one has a free $\mathcal{A}^{\text{st}, \text{Gal}}$ -module*

$$G_{\Phi}^{\text{wk}}$$

of rank $1 + (d - c)$ equipped with a $\Pi_{T^{\log}}$ -action which fits into an exact sequence of $\mathcal{A}^{\text{st}, \text{Gal}}[\Pi_{T^{\log}}]$ -modules:

$$0 \longrightarrow \Theta_{S^{\log}}^{\text{wk}, \text{et}}(1) \otimes_{W(\mathbf{F}_{p^2})} \mathcal{A}^{\text{st}, \text{Gal}} \longrightarrow G_{\Phi}^{\text{wk}} \longrightarrow \mathcal{A}^{\text{st}, \text{Gal}} \longrightarrow 0$$

Moreover, if one restricts from $\Pi_{T^{\log}}$ to $\Pi_{(S^{\text{st}})^{\log}} \stackrel{\text{def}}{=} \pi_1((S^{\text{st}})^{\log}_K)$, then one has a natural $\Pi_{(S^{\text{st}})^{\log}}$ -invariant augmentation

$$\epsilon_{\text{Gal}} : \mathcal{A}^{\text{st}, \text{Gal}} \rightarrow W(\mathbf{F}_{p^2})$$

such that when one base-changes the above exact sequence by ϵ_{Gal} (and replaces the resulting $W(\mathbf{F}_{p^2})$ on the right by \mathbf{Z}_p), one obtains the Tate module of the canonical weak uniformizing p -divisible group associated to the anabelian Frobenius system of Theorem 2.3 (cf. the discussion preceding Definition 2.5).

Remark. Thus, we see that the point of the crystalline induction just performed is the following: Whereas in the discussion preceding Definition 2.5, we constructed a $\Theta_{S^{\log}}^{\text{wk}, \text{et}}(1)$ -torsor over the set \mathcal{U}^{st} (of Definition 2.5), i.e., a $\Theta_{S^{\log}}^{\text{wk}, \text{et}}(1)$ -torsor in the category of set-theoretic functions on \mathcal{U}^{st} , the theory of the present subsection shows that, in fact, this torsor is defined in a certain category of p -adic analytic functions on \mathcal{U}^{st} . Moreover, this analytic structure on the torsor is even crystalline in the sense that it arises from some analytic torsor in the category \mathcal{MF}^{∇} via (a slight generalization of) the theory of [Falt1], §2. Thus, in summary, the essence of the technique of crystalline induction is the following:

Crystalline induction is a technique that allows one to show that certain p -adic analytic objects with Galois action (e.g., torsors)

defined over a Galois set (e.g., \mathcal{U}^{st}) in the category of set-theoretic functions on the set in fact admit a natural analytic structure which is compatible with the Galois action and crystalline in the sense that it arises from an \mathcal{MF}^∇ -type object.

Note that the technique of proof employed here (cf. especially Lemmas 2.9 and 2.10) is a direct reflection of the philosophy underlying the notion of crystalline induction.

Other examples of crystalline induction include the theory of Chapter X (of the present work), and [Mzk1], Chapter V. In these cases, the objects (over a Galois set) that are dealt with are *representations of the geometric fundamental group of a hyperbolic curve* (as opposed to torsors). Note that the reason for the word “induction” is that this technique is a sort of crystalline (or \mathcal{MF}^∇) analogue of the technique of *induction in the theory of group representations*. Indeed, in the example dealt with here, the point of view of the discussion preceding Definition 2.5 gives rise to objects over S^{st} , i.e., crystalline Galois representations of $\Pi_{(S^{\text{st}})^{\log}}$. The induction of the present subsection then is essentially an *induction relative to the inclusion of groups*

$$\Pi_{(S^{\text{st}})^{\log}} \subseteq \Pi_{T^{\log}}$$

and results in crystalline Galois representations of $\Pi_{T^{\log}}$.

Finally, we would like to prepare for the *crystalline inductions* that we will perform in *Chapter X* as follows: Just as we formed $\mathcal{A}^{\text{st}, \text{Gal}}$ out of $(G^{\text{st}})_{\text{Gal}}^{\mathbf{F}_{p^2}}$, we would like to form an $\mathcal{A}^{\text{st}, \text{Gal}}$ -algebra out of G_Φ^{wk} . First, recall the $\Pi_{T^{\log}}$ -invariant *augmentation*

$$\alpha_{\text{Gal}} : \text{Aff}((G^{\text{st}})_{\text{Gal}}^{\mathbf{F}_{p^2}}) \rightarrow \mathbf{F}_{p^2}$$

If we take the dual of the exact sequence of Theorem 2.11, and then base-change by this augmentation, we get an exact sequence

$$0 \rightarrow \mathbf{F}_{p^2} \rightarrow (G_\Phi^{\text{wk}})_{\mathbf{F}_{p^2}}^\vee \rightarrow \Omega_{S_k^{\log}}^{\text{wk}, \text{et}}(-1) \rightarrow 0$$

The covering parametrizing *splittings* of this sequence is a finite, flat morphism

$$Z_K^{\log} \rightarrow T_K^{\log}$$

of degree $p^{2(d-c)}$. Note that Z_K^{\log} is *étale* over T_K^{\log} . Let

$$\Pi_{Z^{\log}} \stackrel{\text{def}}{=} \pi_1(Z_K^{\log})$$

Then if we restrict the action of $\Pi_{T^{\log}}$ on

$$\mathrm{Aff}((G_{\Phi}^{\mathrm{wk}})^{\vee})$$

(which will be a p -adic polynomial algebra of dimension $d - c$ over $\mathcal{A}^{\mathrm{st}, \mathrm{Gal}}$) to $\Pi_{Z^{\log}}$, there exists a tautological $\Pi_{Z^{\log}}$ -equivariant augmentation

$$\mathrm{Aff}((G_{\Phi}^{\mathrm{wk}})^{\vee}) \rightarrow \mathbf{F}_{p^2}$$

Taking the p -adic completion of the PD-envelope at this augmentation, we thus obtain an $\mathcal{A}^{\mathrm{st}, \mathrm{Gal}}[\Pi_{Z^{\log}}]$ -algebra

$$\mathcal{A}_{\Phi}^{\mathrm{Gal}}$$

which will be (non-canonically isomorphic to) the p -adic completion of the PD-envelope (at the origin) of a polynomial algebra of dimension d over $W(\mathbf{F}_{p^2})$.

Definition 2.12. Let us call $\mathcal{A}_{\Phi}^{\mathrm{Gal}}$ (respectively, $\mathcal{A}^{\mathrm{st}, \mathrm{Gal}}$) the *total* (respectively, *strong*) *Galois mantle* associated to $\Phi^{\log} : S^{\log} \rightarrow S^{\log}$ and H .

Now, the argument used above to construct the *key morphism* $(\mathcal{A}_T^{\mathrm{st}})(B^+(T^{\log})) \rightarrow \mathcal{A}_{B^+}^{\mathrm{st}, \mathrm{Gal}}$ for the strong mantle may be repeated for the *total mantle* as follows: First, note that from the definition of G_{Φ}^{wk} , we have a natural $\mathcal{A}^{\mathrm{st}, \mathrm{Gal}}$ -linear morphism

$$G_{\Phi}^{\mathrm{wk}} \stackrel{\mathrm{def}}{=} F^0(\mathcal{M})^{F=1} \subseteq \mathcal{M} = \mathrm{Ind}^{\mathrm{Gal}}(\mathcal{A}) \oplus \mathrm{Ind}^{\mathrm{Gal}}(\mathcal{B}) \rightarrow \mathrm{Ind}^{\mathrm{Gal}}(\mathcal{A})$$

(where the morphism on the right is the natural projection). If we tensor this morphism (on the left) over $\mathcal{A}^{\mathrm{st}, \mathrm{Gal}}$ with $\mathcal{A}_{B^+}^{\mathrm{st}, \mathrm{Gal}}$, we get a $\mathcal{A}_{B^+}^{\mathrm{st}, \mathrm{Gal}}$ -linear morphism of free $\mathcal{A}_{B^+}^{\mathrm{st}, \mathrm{Gal}}$ -modules of finite rank:

$$G_{\Phi}^{\mathrm{wk}} \otimes_{\mathcal{A}^{\mathrm{st}, \mathrm{Gal}}} \mathcal{A}_{B^+}^{\mathrm{st}, \mathrm{Gal}} \rightarrow \mathrm{Ind}^{\mathrm{Gal}}(\mathcal{A})$$

Taking the $\mathcal{A}_{B^+}^{\mathrm{st}, \mathrm{Gal}}$ -linear dual of this morphism, and then applying the definitions of $\mathrm{Ind}^{\mathrm{Gal}}(\mathcal{A})$ and $\mathcal{A}_{\Phi}^{\mathrm{Gal}}$ thus gives a morphism

$$\mathrm{Ind}(\mathcal{A}_T)(B^+(T^{\log})) \rightarrow \mathrm{Ind}^{\mathrm{Gal}}(\mathcal{A})^{\vee} \rightarrow (G_{\Phi}^{\mathrm{wk}})^{\vee} \otimes_{\mathcal{A}^{\mathrm{st}, \mathrm{Gal}}} \mathcal{A}_{B^+}^{\mathrm{st}, \mathrm{Gal}} \hookrightarrow \mathcal{A}_{\Phi}^{\mathrm{Gal}} \otimes_{\mathcal{A}^{\mathrm{st}, \mathrm{Gal}}} \mathcal{A}_{B^+}^{\mathrm{st}, \mathrm{Gal}}$$

On the other hand, since $\mathcal{A}(B^+(S^{\log})) = \mathcal{A}_T(B^+(T^{\log}))$ is the p -adic completion of a certain PD-envelope of the affinization of $\mathrm{Ind}(\mathcal{A}_T)$, we thus obtain a natural morphism

$$\mathcal{A}(B^+(S^{\log})) \rightarrow (\mathcal{A}_{\Phi}^{\mathrm{Gal}})_{B^+} \stackrel{\mathrm{def}}{=} \mathcal{A}_{\Phi}^{\mathrm{Gal}} \otimes_{\mathcal{A}^{\mathrm{st}, \mathrm{Gal}}} \mathcal{A}_{B^+}^{\mathrm{st}, \mathrm{Gal}} = \mathcal{A}_{\Phi}^{\mathrm{Gal}} \otimes_{W(\mathbf{F}_{p^2})} B^+(S^{\log})$$

(where the tensor products are topological, with respect to the p -adic topology). Moreover, by taking the Frobenius pull-back of this morphism, we get a morphism

$$\mathcal{F}'(B^+(S^{\log})) \rightarrow (\mathcal{A}_{\Phi}^{\mathrm{Gal}})_{B^+}^F$$

Note that these two morphisms *respect Hodge filtrations, Frobenius, and the action of $\Pi_{Z^{\log}}$* .

Base-change by these two morphisms

$$\mathcal{A}(B^+(S^{\log})) \rightarrow (\mathcal{A}_{\Phi}^{\mathrm{Gal}})_{B^+}$$

$$\mathcal{F}'(B^+(S^{\log})) \rightarrow (\mathcal{A}_{\Phi}^{\mathrm{Gal}})_{B^+}^F$$

will play an analogous role in the crystalline inductions of Chapter X to the role played by the morphism $(\mathcal{A}_T^{\mathrm{st}})(B^+(T^{\log})) \rightarrow \mathcal{A}_{B^+}^{\mathrm{st}, \mathrm{Gal}}$ in the present subsection.

§2.4. Discussion of the Resulting Spiked Geometry

Just as in Chapter VIII, §2.5, 2.6, where we paused to discuss the meaning of *Lubin-Tate* and *anabelian geometries*, we now pause to discuss the meaning of *spiked geometries* as developed thus far in this Chapter. We begin by making the following remark: Although it is unfortunate that we lack the patience and technical skill to work out the geometry associated to an arbitrary system of Frobenius liftings as in Theorem 1.8 of Chapter VII, we believe that the examples that we have worked out (Lubin-Tate, anabelian, (very ordinary) spiked) are fairly representative of all the main phenomena involved. The reader who is interested may then work out what sorts of geometries arise from more general types of systems of Frobenius liftings. In fact, in some sense, perhaps the present spiked case is the most interesting in that it, by itself, essentially illustrates all the phenomena involved in both the Lubin-Tate and anabelian cases.

Indeed, the *spiked geometry* admits the following overall description: First, one has on S^{\log} a collection of *strong variables* which define a *Lubin-Tate geometry of order 2* on some *virtual (fictitious) quotient*

$$S^{\log} \rightarrow S_{\text{st}}^{\log}$$

of S (Theorem 1.5). The fibers of this quotient are then equipped with an *anabelian geometry of length 2 and order 1* (Theorem 2.3). That is to say:

The essence of a very ordinary spiked geometry (of period 2) is that it is a family of anabelian, length 2, order 1 geometries parametrized by a Lubin-Tate geometry of order 2.

Thus, Theorem 2.11 tells us that this geometry can be completely translated into a geometry governing the Galois action on the Galois mantle. We give an illustration of the Galois mantle in the Pictorial Appendix.

§2.5. Construction of the Galois Mantle: The Binary-Ordinary Case

While we are thinking about Galois mantles, we continue our preparation for the *crystalline inductions of Chapter X* by making the following observation: It is essentially trivial from the theory of Chapter VIII to form Galois mantles associated to an arbitrary binary-ordinary Frobenius system. Indeed, let us assume that we are given a *binary-ordinary system of Frobenius liftings*

$$\Phi_i^{\log} : S^{\log} \rightarrow S^{\log}$$

(where $i = 1, \dots, n$) of multi-order $\Lambda = \{\lambda_1, \dots, \lambda_n\}$ and period ϖ (cf. Definitions 1.1, 2.1 of Chapter VIII). For $i = 1, \dots, n$, write

$$(\mathcal{A}[i], \nabla_{\mathcal{A}[i]}, F^*(\mathcal{A}[i])) \stackrel{\text{def}}{=} (\mathcal{A}, \nabla_{\mathcal{A}}, F^*(\mathcal{A}))$$

Then Φ_i^{\log} defines a morphism

$$(\Phi_{S_{\mathbb{F}_p}}^*)^{\lambda_i-1} \mathbf{F}^*(\mathcal{A}[i], \nabla_{\mathcal{A}[i]}) \rightarrow (\mathcal{A}[i+1], \nabla_{\mathcal{A}[i+1]})$$

(where \mathbf{F}^* is the *renormalized Frobenius pull-back of the mantle*, as defined in Definition 1.10) which becomes an isomorphism if one takes the p -adic completion of the PD-envelope of the left-hand side at the section (of the left-hand side) defined (cf. Lemma 1.9) by Φ_i^{\log} .

Now let us look at *Galois representations*, i.e., representations of $\Pi_{S^{\log}}$. Recall the exact sequence of (log) p -divisible groups

$$0 \rightarrow G_{\Omega_{\Phi}} \rightarrow G_{\Phi} \rightarrow \mathbf{Q}_p/\mathbf{Z}_p \rightarrow 0$$

in the discussion preceding Definition 2.9 of Chapter VIII. If we apply $\mathrm{Hom}(-, \mathbf{Q}_p/\mathbf{Z}_p)$ to this sequence, push-forward the resulting \mathbf{Z}_p on the left by $\mathbf{Z}_p \hookrightarrow \mathcal{O}_\varpi$, and restrict to S_K^{log} , we obtain a sequence of $\Pi_{S^{\mathrm{log}}}$ -modules

$$0 \rightarrow \mathcal{O}_\varpi \rightarrow G_\Phi^{\mathbf{F}_{p^\varpi}} \rightarrow T_p(G_{\Omega_\Phi}) \rightarrow 0$$

Let

$$Z_K^{\mathrm{log}} \rightarrow S_K^{\mathrm{log}}$$

be the covering (of degree $p^{\varpi \cdot d}$) parametrizing *splittings* of this sequence modulo p . Then after restricting to $\Pi_{Z^{\mathrm{log}}}$, $\mathrm{Aff}(G_\Phi^{\mathbf{F}_{p^\varpi}})$ admits a tautological $\Pi_{Z^{\mathrm{log}}}$ -invariant augmentation

$$\mathrm{Aff}(G_\Phi^{\mathbf{F}_{p^\varpi}}) \rightarrow \mathbf{F}_{p^\varpi}$$

Taking the p -adic completion of the PD-envelope at this augmentation, we thus obtain a $\mathcal{O}_\varpi[\Pi_{Z^{\mathrm{log}}}]$ -algebra

$$\mathcal{A}_\Phi^{\mathrm{Gal}}$$

which will be (non-canonically isomorphic to) the p -adic completion of the PD-envelope (at the origin) of a polynomial algebra of dimension d over $W(\mathbf{F}_{p^\varpi})$.

Definition 2.13. Let us call $\mathcal{A}_\Phi^{\mathrm{Gal}}$ the *Galois mantle* associated to the binary-ordinary Frobenius system $\{\Phi_i^{\mathrm{log}}\}_{i=1, \dots, n}$.

If $i \in \{1, \dots, n\}$, and $j \in \{0, \dots, \lambda_i - 1\}$, then let

$$(\mathcal{A}[i, j], \nabla_{\mathcal{A}[i, j]}) \stackrel{\mathrm{def}}{=} (\Phi_{S_{\mathbf{F}_p}}^*)^{j-1} \mathbf{F}^*(\mathcal{A}[i], \nabla_{\mathcal{A}[i]})$$

(if $j > 0$) and

$$(\mathcal{A}[i, j], \nabla_{\mathcal{A}[i, j]}) \stackrel{\mathrm{def}}{=} (\mathcal{A}[i], \nabla_{\mathcal{A}[i]})$$

(if $j = 0$). Then, by the exact same argument as that used to construct the *key morphism* $(\mathcal{A}_T^{\mathrm{st}})(B^+(T^{\mathrm{log}})) \rightarrow \mathcal{A}_{B^+}^{\mathrm{st}, \mathrm{Gal}}$ of §2.3, we may construct *morphisms*

$$(\mathcal{A}[i, j])(B^+(S^{\mathrm{log}})) \rightarrow (\mathcal{A}_\Phi^{\mathrm{Gal}})_{B^+}^{F^{N_{i,j}}}$$

(where $(\mathcal{A}_{\Phi}^{\text{Gal}})_{B^+} \stackrel{\text{def}}{=} B^+(S^{\log}) \otimes_{\mathcal{O}_{\varpi}} \mathcal{A}_{\Phi}^{\text{Gal}}$, and $N_{i,j} = j + \sum_{a=1}^{i-1} \lambda_a$) that respect Hodge filtrations, Frobenius, and the action of $\Pi_{Z^{\log}}$. Moreover,

Base-change by these morphisms will play a crucial role in the crystalline inductions that we will carry out in Chapter X (cf. what we did in §2.3).

§3. Application to Curves and their Moduli

In this §, we briefly note how the general theory of the preceding two §'s can be applied in the case of hyperbolic curves and their moduli stacks.

§3.1. Frobenius Liftings on the Moduli Stack

Let g, r be nonnegative integers such that $2g - 2 + r \geq 1$. Let Π be a *spiked VF-pattern of period* $\varpi = 2$. Let

$$S^{\log} \stackrel{\text{def}}{=} ((\overline{\mathcal{N}}_{g,r}^{\Pi,s})^{\log})_{\mathbf{Z}_p}^{\text{very ord}}$$

be the *very* Π -ordinary locus (cf. Definition 3.6 of Chapter VII) of $((\overline{\mathcal{N}}_{g,r}^{\Pi,s})^{\log})_{\mathbf{Z}_p}^{\text{ord}}$. Thus, S^{\log} is a p -adic formal log stack whose underlying stack is formally étale over $(\overline{\mathcal{M}}_{g,r})_{\mathbf{Z}_p}$ and formally smooth over \mathbf{Z}_p . Let

$$X^{\log} \rightarrow S^{\log}$$

be the *tautological curve* over S^{\log} . According to Theorem 1.8 of Chapter VII, we thus have a *canonical Frobenius lifting*

$$\Phi^{\log} : S^{\log} \rightarrow S^{\log}$$

In this subsection, we would like to show that Φ^{\log} is a *very ordinary spiked Frobenius lifting of colevel* $2(\chi - \Pi(1))$ (in the sense of Definition 1.1) and, moreover, admits a natural *Hodge subspace* (cf. Definition 1.13).

We begin with Definition 1.1. Note that condition (1) follows immediately from Chapter VII, Theorem 3.8: Indeed, (in the notation of *loc. cit.*) the vector bundle $\Omega_{S_{\mathbf{F}_p}}^{\log}$ is given by the subbundle

$$\theta_{\mathbf{F}_p}^* \Omega'_{T_{\mathbf{F}_p}^{\log}} \subseteq \theta_{\mathbf{F}_p}^* \Omega_{T_{\mathbf{F}_p}^{\log}} \cong (\Phi_{S_{\mathbf{F}_p}}^2)^* \Omega_{S_{\mathbf{F}_p}}^{\log}$$

Note that the “ T ” of Chapter VII, §3.2, §3.3, is *not* the same as the “ T ” of §1.2, 2.2, of the present Chapter. Indeed, roughly speaking, the “ $T_{\mathbf{F}_p}$ ” of the present Chapter is obtained by taking p^2 -th roots of the strong variables, whereas the “ $T_{\mathbf{F}_p}$ ” of Chapter VII, §3.2, 3.3, is obtained by taking p -th roots of the strong variables. Condition (3) of Definition 1.1 then follows immediately from the isomorphisms discussed at the end of the statement of Chapter VII, Theorem 3.8. Condition (2) of Definition 1.1 follows from the fact that Φ^{\log} can be factored as a *product of two Frobenius liftings* (cf. Chapter VII, Theorem 3.8) both of which carry divisors at infinity to p times themselves.

Finally, we consider Condition (4) of Definition 1.1. First, relative to the discussion of Chapter VII, §3.2, observe that the superscript “Hs” (of *loc. cit.*) (respectively, “spk”) corresponds – in the language of the present Chapter – to the “weak variables” (respectively, “strong variables”). That is, we have:

$$\text{Hs} \iff \text{weak variables} \qquad \text{spk} \iff \text{strong variables}$$

First, since the discussion of Chapter VII, §3.2, is from the point of view of *tangent bundles* (as opposed to bundles of differentials), let us reword Condition (4) in terms of tangent bundles. In the notation of *loc. cit.*, we have morphisms

$$\Theta_{S_{\mathbf{F}_p}^{\log}} \xrightarrow{\delta^{\text{Hs}}} (\Phi^{\text{Hs}})^* \Theta_{T_{\mathbf{F}_p}^{\log}} \quad (\Phi^{\text{Hs}})^* \Theta_{\kappa} \xrightarrow{(\Phi^{\text{Hs}})^* \Theta_{\kappa}} (\Phi_{S_{\mathbf{F}_p}^2})^* \Theta_{S_{\mathbf{F}_p}^{\log}}$$

where $\delta^{\text{Hs}} = (\text{d}\theta_{\mathbf{F}_p})^{\vee}$ and $(\Phi^{\text{Hs}})^* \Theta_{\kappa} = (p^{-1} \cdot \text{d}\phi)^{\vee}$. Moreover, δ^{Hs} factors as

$$\Theta_{S_{\mathbf{F}_p}^{\log}} \rightarrow (\Phi^{\text{Hs}})^* \Theta_T^{\text{spk}} \hookrightarrow (\Phi^{\text{Hs}})^* \Theta_{T_{\mathbf{F}_p}^{\log}}$$

In the following, let us denote the surjection on the left by $\pi^{\text{spk}} : \Theta_{S_{\mathbf{F}_p}^{\log}} \rightarrow (\Phi^{\text{Hs}})^* \Theta_T^{\text{spk}}$. Thus, the composite $\{(\Phi^{\text{Hs}})^* \Theta_{\kappa}\} \circ \delta^{\text{Hs}}$ induces a morphism

$$(\Phi^{\text{Hs}})^* \Theta_T^{\text{spk}} \rightarrow (\Phi_{S_{\mathbf{F}_p}^2})^* \Theta_{S_{\mathbf{F}_p}^{\log}}$$

which we may compose with the pull-back by $\Phi_{S_{\mathbf{F}_p}^2}^2$ of the projection π^{spk} to obtain a morphism

$$\tau^{\text{spk}} : (\Phi^{\text{Hs}})^* \Theta_T^{\text{spk}} \rightarrow (\Phi_{S_{\mathbf{F}_p}^2})^* (\Phi^{\text{Hs}})^* \Theta_T^{\text{spk}}$$

The content of Condition (4) is that this morphism τ^{spk} be an isomorphism. This is clearly equivalent to the statement that

$$(\Phi^{\text{Hs}})^* \Theta_{T_{\mathbf{F}_p}^{\log}} = \{(\Phi^{\text{Hs}})^* \Theta_T^{\text{spk}}\} \oplus (\Phi^{\text{Hs}})^* \Theta_{\kappa}^{-1} \{(\Phi_{S_{\mathbf{F}_p}}^* \text{Ker}(\pi^{\text{spk}}))|_T\}$$

which, in turn, is equivalent (since Φ^{Hs} is *faithfully flat*, and $\text{Ker}(\pi^{\text{spk}})|_T = \text{Ker}(\delta^{\text{Hs}})|_T = \Theta_T^{\text{Hs}}$ – cf. Chapter VII, Lemma 3.3) to the statement

$$\Theta_{T_{\mathbf{F}_p}^{\log}} = (\Theta_T^{\text{spk}}) \oplus \Theta_{\kappa}^{-1} (\Phi_{T_{\mathbf{F}_p}}^* \Theta_T^{\text{Hs}})$$

(where we regard Θ_T^{Hs} as a subbundle of $\Theta_{S|T}$). On the other hand, since $\Theta_{\mathbf{Q}_1}$ maps Θ_T^{spk} onto $F^0/F^1((\Theta_1)_T)$, it thus follows that Condition (4) is equivalent to the statement that the morphism

$$\Phi_{T_{\mathbf{F}_p}}^* \Theta_T^{\text{Hs}} \rightarrow F^{-1}/F^0((\Theta_1)_T)$$

(considered in (b.) of Chapter VII, Definition 3.6!) is an isomorphism. Thus, *Condition (4) of Definition 1.1 of the present Chapter is a consequence of (b.) of Chapter VII, Definition 3.6.*

Next, we construct a *Hodge subspace* $H \subseteq F'$ as in Definition 1.13. (Here,

$$F' = \text{Spf}(\mathbf{F}^*(\mathcal{A}, \nabla_{\mathcal{A}}))$$

(cf. Definition 1.10) where $(\mathcal{A}, \nabla_{\mathcal{A}})$ is the mantle of S^{\log} .) To do this, recall the morphism

$$\gamma' : F' \rightarrow S$$

(from the discussion preceding Lemma 1.9) obtained from the inclusion $\mathcal{O}_S \hookrightarrow \mathcal{A}$ from the *left*. Let

$$\mathcal{P} = (P, \nabla_P)$$

be the *canonical indigenous bundle* (Theorem 1.8 of Chapter VII) on X^{\log} . Let

$$X_{\gamma'}^{\log}; \quad \mathcal{P}_{\gamma'}$$

be the pull-backs of X^{\log} and \mathcal{P} via γ' . Note that $(X_{\gamma'}^{\log})_{\mathbf{F}_p} \rightarrow (F')_{\mathbf{F}_p}^{\log}$ is simply the result of base-changing $(X_{\mathbf{F}_p}^{\log})^F \rightarrow S_{\mathbf{F}_p}^{\log}$ via the *structure*

morphism $\beta' : F' \rightarrow S^{\log}$. Thus, we can form the *renormalized Frobenius pull-back* of $\mathcal{P}_{\gamma'}$ to obtain a crystal in \mathbf{P}^1 -bundles

$$\mathcal{Q}$$

on $\text{Crys}((X_{\beta'}^{\log})_{\mathbf{F}_p}/F')$, where $X_{\beta'}^{\log} \rightarrow F'$ is the result of pulling back X^{\log} via the structure morphism $\beta' : F' \rightarrow S$. Now we make the following definition:

We define $H' \subseteq F' \rightarrow S$ to be the S -subscheme (cf. Chapter I, Theorem 3.10) where \mathcal{Q} defines a *crys-stable bundle of level $\Pi(1)$* on $X_{\beta'}^{\log}$.

Our *Hodge subspace* $H \subseteq F'$ will be a certain *open subscheme* of H' (to be specified below).

Let us continue to write $T_{\mathbf{F}_p}$ for the object “ T ” of Chapter VII, §3.2. Note that the Frobenius lifting $\phi_{\mathbf{Z}_p}^{\log} : T_{\mathbf{Z}_p}^{\log} \rightarrow S_{\mathbf{Z}_p}^{\log}$ of Chapter VII, Theorem 3.8, thus defines (by the functorial interpretation of $F'_{\mathbf{F}_p}$ – cf. Lemma 1.9) a *section*

$$\xi_T : T_{\mathbf{F}_p} \rightarrow F'|_{T_{\mathbf{F}_p}}$$

Alternatively, ξ_T may be regarded as the section defined by the \mathbf{P}^1 -bundle \mathbf{Q}_1 of Chapter VII, §3.2. Thus, the Kodaira-Spencer morphism

$$\Theta_{T_{\mathbf{F}_p}^{\log}} \rightarrow \Phi_{T_{\mathbf{F}_p}}^*(\Theta_{S_{\mathbf{F}_p}^{\log}})|_T = \xi_T^* \Theta_{F'/S}$$

of this section ξ_T may be identified with the morphism Θ_{κ} of Chapter VII, §3.2, hence (by Chapter VII, Definition 3.6, (a.)) is an *isomorphism*. Note that since \mathbf{Q}_1 is a *crys-stable bundle of level $\Pi(1)$* , we also conclude that *the image of ξ_T lies in H'* . On the other hand, it follows from the definition of the morphism ξ (appearing in Definition 1.13) and the fact that the Kodaira-Spencer morphism of ξ_T is an *isomorphism* that the image of ξ is equal to the image of ξ_T , hence isomorphic to $T_{\mathbf{F}_p}$. Since $T_{\mathbf{F}_p}$ is k -smooth, this shows that Condition (1) of Definition 1.13 is satisfied.

Next, let us observe that since the morphism $\Theta_{\mathbf{Q}_1}$ (of Chapter VII, §3.2) maps Θ_T^{spk} onto $F^0/F^1((\Theta_1)_T)$, it follows (from the definition of H' – cf. also Chapter I, Lemma 3.8) that the Kodaira-Spencer morphism $\Theta_{T_{\mathbf{F}_p}^{\log}} \cong \xi_T^* \Theta_{F'/S}$ maps $\Theta_T^{\text{spk}} \subseteq \Theta_{T_{\mathbf{F}_p}^{\log}}$ isomorphically onto $\xi_T^* \Theta_{H'/S} \subseteq \xi_T^* \Theta_{F'/S}$. Thus, the morphism

$$\xi_T^* \Theta_{H'/S} \rightarrow \xi_T^* \Theta_{(F')^{\text{st}}/S}$$

on tangent bundles arising from the projection considered in Condition (2) of Definition 1.13 may (after further pull-back by Φ^{Hs}) be identified with the morphism

$$\tau^{\text{spk}} : (\Phi^{\text{Hs}})^* \Theta_T^{\text{spk}} \rightarrow (\Phi_{S_{\mathbb{F}_p}}^2)^* (\Phi^{\text{Hs}})^* \Theta_T^{\text{spk}}$$

considered above. Since we have already seen that τ^{spk} is an isomorphism, we thus conclude that $H' \rightarrow (F')^{\text{st}}$ is étale in a neighborhood of the image of ξ_T , i.e., Condition (2) of Definition 1.13 is satisfied in a neighborhood of the image of ξ_T .

Finally, we consider Condition (3) of Definition 1.13. Note that we have a commutative diagram of exact sequences

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \Theta_T^{\text{spk}} & \longrightarrow & \Theta_{T_{\mathbb{F}_p}^{\log}} & \longrightarrow & \Theta_T^{\text{Hs}} & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & \xi_T^* \Theta_{H'/S} & \longrightarrow & \xi_T^* \Theta_{F'/S} & \longrightarrow & \xi_T^* \mathcal{N}_{H'/F'} & \longrightarrow & 0 \end{array}$$

induced by the Kodaira-Spencer morphism of ξ_T in which all the vertical morphisms are isomorphisms. On the other hand, the morphism

$$\delta^{\text{spk}} : \Theta_{T_{\mathbb{F}_p}^{\log}} \rightarrow \Theta_{S_{\mathbb{F}_p}^{\log}}|_T$$

maps the quotient $\Theta_{T_{\mathbb{F}_p}^{\log}} \rightarrow \Theta_T^{\text{Hs}}$ isomorphically onto

$$\Theta_{S_{\mathbb{F}_p}^{\log}}|_T = (\Omega_{S_{\mathbb{F}_p}^{\log}}^{\text{wk}})^{\vee}|_T \subseteq \Omega_{S_{\mathbb{F}_p}^{\log}}^{\vee}|_T = \Theta_{S_{\mathbb{F}_p}^{\log}}|_T$$

(cf. Lemma 2.2). Thus, the fact that the vertical morphism on the right (of the above commutative diagram of exact sequences) is an isomorphism implies that Condition (3) of Definition 1.13 is satisfied in a neighborhood of the image of ξ_T .

Thus, by taking our Hodge subspace $H \subseteq H'$ to be a suitable neighborhood of the image of ξ_T , we see that all the conditions of Definition 1.13 are satisfied. In other words, we have proven the following result:

Theorem 3.1. *Let*

$$S^{\log} \stackrel{\text{def}}{=} ((\overline{\mathcal{N}}_{g,r}^{\Pi,s})^{\log})_{\mathbb{Z}_p}^{\text{very ord}}$$

be the very Π -ordinary locus (Definition 3.6 of Chapter VII) of $((\overline{\mathcal{N}}_{g,r}^{\Pi,s})^{\log})_{\mathbb{Z}_p}^{\text{ord}}$, where $2g - 2 + r > 0$ and Π is spiked of period 2. Then the associated canonical Frobenius

lifting $\Phi^{\log} : S^{\log} \rightarrow S^{\log}$ of Chapter VII, Theorem 1.8, is a very ordinary spiked Frobenius lifting of colevel $2(\chi - \Pi(1))$ (in the sense of Definition 1.1), and, moreover, there exists a natural choice of Hodge subspace $H \subseteq F'$ (obtained as the locus of crys-stable bundles of level $\Pi(1)$) for this Frobenius lifting.

In particular, one obtains on S^{\log} (and hence, étale locally on $(\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{Z}_p}$) strong uniformizations as in Theorem 1.5, weak anabelian uniformizations as in Theorem 2.3, and a uniformizing Galois representation as in Theorem 2.11.

Remark. In particular, when $(g, r) = (0, 5)$, we obtain a purely Lubin-Tate geometry of order 2 on S . This is somewhat surprising in that since there are marked points, the curve itself can never admit such a geometry.

§3.2. Frobenius Liftings on the Universal Curve

We maintain the notation of §3.1. Let us define the *ordinary locus* of X^{\log}

$$(X^{\log})_{\mathbf{F}_p}^{\text{ord}} \subseteq X_{\mathbf{F}_p}^{\log}$$

to be the open set over which the Kodaira-Spencer morphisms of the Hodge sections of \mathbf{P}_0 and \mathbf{Q}_1 , as well as the Hasse-invariant type morphisms

$$\omega_{X_{\mathbf{F}_p}^{\log}/S_{\mathbf{F}_p}^{\log}} \rightarrow \text{Ad}(\mathbf{P}_0) \rightarrow \Phi_{X_{\mathbf{F}_p}}^* \omega_{X_{\mathbf{F}_p}^{\log}/S_{\mathbf{F}_p}^{\log}}$$

$$\omega_{X_{\mathbf{F}_p}^{\log}/S_{\mathbf{F}_p}^{\log}}|_T \rightarrow \text{Ad}(\mathbf{Q}_1) \rightarrow (\Phi_{X_{\mathbf{F}_p}}^* \omega_{X_{\mathbf{F}_p}^{\log}/S_{\mathbf{F}_p}^{\log}})|_T$$

(where \mathbf{P}_0 , \mathbf{Q}_1 , and T are as in Chapter VII, §3.2) are isomorphisms. (Here, the morphisms on the left are derived from the Hodge filtration and its Kodaira-Spencer map, while the morphisms on the right are the duals of the p -curvature for \mathbf{P}_0 and \mathbf{Q}_1 , respectively.) Let

$$(X^{\log})^{\text{ord}} \subseteq X^{\log}$$

be the p -adic formal open substack defined by $(X^{\log})_{\mathbf{F}_p}^{\text{ord}}$. Then we would like to define (in this subsection) a *very ordinary spiked Frobenius lifting* on $(X^{\log})^{\text{ord}}$ that covers $\Phi^{\log} : S^{\log} \rightarrow S^{\log}$, together with a *Hodge subspace* $H_X \subseteq F'_X$. (Here, F'_X denotes the renormalized Frobenius pull-back of the mantle on $(X^{\log})^{\text{ord}}$.)

First, we construct the *Frobenius lifting* (étale locally) on $(X^{\log})^{\text{ord}}$. Let us denote by

$$X_{\Phi}^{\log} \stackrel{\text{def}}{=} X^{\log} \times_{S^{\log}, \Phi^{\log}} S^{\log} \rightarrow S^{\log}$$

$$\mathcal{P}_{\Phi} \stackrel{\text{def}}{=} \mathcal{P}|_{X_{\Phi}^{\log}}$$

(where, just as in §3.1, $\mathcal{P} = (P, \nabla_P)$ denotes the *canonical indigenous bundle* on X^{\log} – cf. Chapter VII, Theorem 1.8). Note that, since the *naive and renormalized Frobenius pull-backs* become the same after one tensors with \mathbf{Q}_p , we have:

$$(\Phi_{X_{\mathbf{F}_p}/S_{\mathbf{F}_p}}^*)^2 \mathcal{P}_{\Phi} \otimes \mathbf{Q}_p \cong \mathcal{P} \otimes \mathbf{Q}_p$$

If we restrict this isomorphism to $(X^{\log})^{\text{ord}}$, then *any lifting*

$$\Psi^{\log} : (X^{\log})^{\text{ord}} \rightarrow (X^{\log})^{\text{ord}}$$

of the square of Frobenius that covers $\Phi^{\log} : S^{\log} \rightarrow S^{\log}$ *determines an isomorphism*

$$\zeta_{\Psi} : (\Psi^* P)_{\mathbf{Q}_p}|_{(X^{\log})^{\text{ord}}} \cong P_{\mathbf{Q}_p}|_{(X^{\log})^{\text{ord}}}$$

Now we have the following:

Lemma 3.2. *If there exists a Ψ^{\log} as above such that ζ_{Ψ} is compatible with the Hodge sections of P on either side, then this Ψ^{\log} is unique.*

Proof. Let

$$W \stackrel{\text{def}}{=} W((S_{\mathbf{F}_p}^{\log})^{\text{pro}})$$

be the Witt scheme associated to the perfection of $S_{\mathbf{F}_p}^{\log}$. Then Φ^{\log} defines a morphism

$$W \rightarrow S^{\log}$$

(cf. Chapter VIII, §1.1), hence an inclusion on points $(X^{\log})^{\text{ord}}(S^{\log}) \hookrightarrow (X^{\log})^{\text{ord}}(W)$. Let us choose a point

$$x \in (X^{\log})^{\text{ord}}(S^{\log})$$

that avoids the marked points and nodes of X . (Indeed, this is always possible étale locally on S . But Lemma 3.2 is true if and only if it is true étale locally.) Let us denote by

$$X_x^{\text{PD}}$$

the p -adic completion of the PD-envelope of $X^{\log} \times_{S^{\log}} W$ (over W) at the point

$$x \in (X^{\log})^{\text{ord}}(S^{\log}) \hookrightarrow (X^{\log})^{\text{ord}}(W)$$

Since X_x^{PD} only depends on the reduction of x modulo p , it makes sense to speak of the square Frobenius conjugate

$$X_{x^{F^2}}^{\text{PD}}$$

of X_x^{PD} (i.e., the p -adic completion of the PD-envelope at the square Frobenius conjugate of x modulo p). Let

$$\mathbf{P}_W^1 \rightarrow W$$

be the restriction of $P \rightarrow X$ defined by x . Since (P, ∇_P) is (i.e., defines) a crystal on $\text{Crys}((X^{\log} \times W)_{\mathbf{F}_p}/W)$, \mathbf{P}_W^1 is determined by the reduction of x modulo p ; thus, we also have $(\mathbf{P}_W^1)^{F^2}$. Note that \mathbf{P}_W^1 is equipped with a Frobenius action

$$\Phi_{\mathbf{P}_W^1} : (\mathbf{P}_W^1)_{\mathbf{Q}_p} \stackrel{\text{def}}{=} (\mathbf{P}_W^1) \otimes \mathbf{Q}_p \rightarrow (\mathbf{P}_W^1)^{F^2}$$

which is an isomorphism over \mathbf{Q}_p (i.e., after one also tensors $(\mathbf{P}_W^1)^{F^2}$ with \mathbf{Q}_p).

Now by integrating (P, ∇_P) over X_x^{PD} , we obtain a local trivialization of (P, ∇_P) over X_x^{PD} , i.e., a horizontal isomorphism

$$P|_{X_x^{\text{PD}}} \cong \mathbf{P}_W^1 \times_W X_x^{\text{PD}}$$

Thus, by composing this isomorphism with the Hodge section of $P \rightarrow X$, we obtain a morphism

$$h : X_x^{\text{PD}} \rightarrow \mathbf{P}_W^1$$

which is “PD-étale” (i.e., induces an isomorphism between PD-neighborhoods of W -valued points at x). (Indeed, this follows from the fact that the Kodaira-Spencer morphism of the Hodge section of $P \rightarrow X$ is an isomorphism – since \mathcal{P} is indigenous.) Of course, we

also have a square Frobenius conjugate version $h^{F^2} : X_{x^{F^2}}^{\text{PD}} \rightarrow (\mathbf{P}_W^1)^{F^2}$ of h .

Let us denote by $(X_x^{\text{PD}})_{\mathbf{Q}_p}$ the formal neighborhood of X_x^{PD} at the $(W \otimes \mathbf{Q}_p)$ -valued point defined by x . Denote by $h_{\mathbf{Q}_p} : (X_x^{\text{PD}})_{\mathbf{Q}_p} \rightarrow (\mathbf{P}_W^1)_{\mathbf{Q}_p}$ the morphism defined by h . Thus, it suffices to prove that any Ψ^{PD} that makes the diagram

$$\begin{array}{ccc} (X_x^{\text{PD}})_{\mathbf{Q}_p} & \xrightarrow{h_{\mathbf{Q}_p}} & (\mathbf{P}_W^1)_{\mathbf{Q}_p} \\ \downarrow \Psi^{\text{PD}} & & \downarrow \Phi_{\mathbf{P}_W^1} \\ X_{x^{F^2}}^{\text{PD}} & \xrightarrow{h^{F^2}} & (\mathbf{P}_W^1)^{F^2} \end{array}$$

commute is *unique*. But this follows immediately from the fact that h , hence also h^{F^2} , is PD-étale. This completes the proof. \square

Now let us write Z^{\log} for the object denoted “ $T_{\mathbf{Z}_p}^{\log}$ ” in Chapter VII, §3.3. Recall the morphisms

$$\theta_{\mathbf{Z}_p}^{\log} : S^{\log} \rightarrow Z^{\log}; \quad \phi_{\mathbf{Z}_p}^{\log} : Z^{\log} \rightarrow S^{\log}$$

of Chapter VII, Theorem 3.8. Recall that (by the discussion of Chapter VII, §3.3) we have a crystal $(\mathbf{Q}_1)_{\mathbf{Z}_p}$ – which we denote here by

$$\mathcal{Q}_1$$

– on $\text{Crys}(X_{\mathbf{F}_p}^{\log}|_{Z_{\mathbf{F}_p}}/Z^{\log})$ whose reduction modulo p is the crys-stable bundle (of level $\Pi(1)$) \mathbf{Q}_1 . Here, by “ $|_{Z_{\mathbf{F}_p}}$,” we mean restriction with respect to the morphism “ $\Phi^{\text{spk}} : T \rightarrow S$ ” of Chapter VII, §3.2. (Note that “ T ” (respectively, “ S ”) of *loc. cit.* corresponds to $Z_{\mathbf{F}_p}$ (respectively, $S_{\mathbf{F}_p}$) in the present discussion.)

Lemma 3.3. *There exists a Ψ^{\log} as above such that ζ_{Ψ} is compatible with the Hodge sections of P on either side.*

Proof. By the uniqueness of Lemma 3.2, it suffices to construct Ψ^{\log} étale locally on S . Thus, we may assume that there exists a curve

$$Y^{\log} \rightarrow Z^{\log}$$

lifting $(X_{\mathbf{F}_p}^{\log})|_{Z_{\mathbf{F}_p}} \rightarrow Z_{\mathbf{F}_p}^{\log}$ such that \mathcal{Q}_1 acquires a Hodge section when evaluated on Y^{\log} . (Note that such a Y^{\log} always exists étale locally on S – cf. Chapter I, Proposition 1.7, Lemma 3.8.) Write

$$(Y^{\log})^{\text{ord}} \subseteq Y^{\log}$$

for the open subobject defined by $(X^{\text{ord}})_{\mathbf{F}_p}|_{Z_{\mathbf{F}_p}} \subseteq X_{\mathbf{F}_p}|_{Z_{\mathbf{F}_p}}$. Similarly, we have $(X^{\log})_{\phi}^{\text{ord}} \subseteq X_{\phi}^{\log}$ (where $X_{\phi}^{\log} \stackrel{\text{def}}{=} X^{\log} \times_{S^{\log}, \phi_{\mathbf{Z}_p}^{\log}} Z^{\log}$). Note that to give a morphism $(Y^{\log})^{\text{ord}} \rightarrow (X^{\log})^{\text{ord}}$ that covers $\phi_{\mathbf{Z}_p}^{\log} : Z^{\log} \rightarrow S^{\log}$ is the same as giving a Z^{\log} -linear morphism $(Y^{\log})^{\text{ord}} \rightarrow (X^{\log})_{\phi}^{\text{ord}}$. Next, observe that

$$\mathcal{Q}_1 = \mathbf{F}^*(\mathcal{P}|_{X_{\phi}^{\log}})$$

Moreover, over the ordinary locus, the Kodaira-Spencer morphism of the Hodge sections of \mathcal{Q}_1 and \mathcal{P}_1 is an *isomorphism*. Thus, by the same deformation argument as that used in Chapter VIII, §3.2, we obtain the *existence* of a unique

$$\Psi_0^{\log} : (Y^{\log})^{\text{ord}} \rightarrow (X^{\log})^{\text{ord}}$$

that covers $\phi_{\mathbf{Z}_p}^{\log}$ and respects the Hodge sections of \mathcal{Q}_1 (over $(Y^{\log})^{\text{ord}}$) and \mathcal{P} (over $(X^{\log})^{\text{ord}}$). Moreover, (just as in Chapter VIII, §3.2) Ψ_0^{\log} is *relatively ordinary* over $\phi_{\mathbf{Z}_p}^{\log} : Z^{\log} \rightarrow S^{\log}$ in the sense that the morphism $p^{-1} \cdot d\Psi_0^{\log}$ on the relative differentials is an *isomorphism*. Similarly, since

$$\mathcal{P} = \mathbf{F}^*(\mathcal{Q}_1|_{Y_{\theta}^{\log}})$$

(where $Y_{\theta}^{\log} \stackrel{\text{def}}{=} Y^{\log} \times_{Z^{\log}, \theta_{\mathbf{Z}_p}^{\log}} S^{\log}$), we conclude that there exists a unique relatively ordinary (over $\theta_{\mathbf{Z}_p}^{\log}$)

$$\Psi_1^{\log} : (X^{\log})^{\text{ord}} \rightarrow (Y^{\log})^{\text{ord}}$$

that covers $\theta_{\mathbf{Z}_p}^{\log}$ and respects the Hodge sections of \mathcal{P} (over $(X^{\log})^{\text{ord}}$) and \mathcal{Q}_1 (over $(Y^{\log})^{\text{ord}}$). Thus, taking $\Psi^{\log} \stackrel{\text{def}}{=} \Psi_0^{\log} \circ \Psi_1^{\log}$ completes the proof. \circ

Let us denote the *unique Frobenius lifting* of Lemma 3.3 by

$$\Phi_X^{\log} : (X^{\log})^{\text{ord}} \rightarrow (X^{\log})^{\text{ord}}$$

Moreover, it is easy to see that Φ_X^{\log} is a *very ordinary spiked Frobenius lifting* (cf. Definition 1.1) of colevel $2(\chi - \Pi(1))$. Indeed, Conditions (1), (3)

and (4) follow from the structure of the objects involved modulo p (cf. Chapter VII, §3.2) and the fact (observed in the proof of Lemma 3.3) that Ψ_0^{\log} and Ψ_1^{\log} are *relatively ordinary* over $\phi_{\mathbf{Z}_p}^{\log}$ and $\theta_{\mathbf{Z}_p}^{\log}$. Condition (2) follows from the fact that Ψ^{\log} (in Lemma 3.3) is constructed as the *composite* of Ψ_0^{\log} and Ψ_1^{\log} .

Remark. Note that the reason that (unlike the situation in Chapter VIII, §3.2) we needed to separate the construction of Φ_X^{\log} into *two parts* (i.e., Lemmas 3.2 and 3.3) is that (unlike in the binary-ordinary case) in the present spiked case, we do not have a *natural choice of lifting* $Y^{\log} \rightarrow Z^{\log}$ (as in the proof of Lemma 3.3) – cf. the difficulties encountered in the construction of $\phi_{\mathbf{Z}_p}^{\log}$ and $\theta_{\mathbf{Z}_p}^{\log}$ in Chapter VII, §3.3.

Now let us construct the *Hodge subspace* H_X . Let

$$\gamma'_X : F'_X \rightarrow X^{\text{ord}}$$

be the morphism “ γ ” for F'_X (derived from the inclusion of $\mathcal{O}_{X^{\text{ord}}}$ into the mantle of X^{ord} from the *left*). (Similarly, we have the *structure morphism* $\beta'_X : F'_X \rightarrow X^{\text{ord}}$.) Note that we have a natural *projection morphism*

$$F'_X \rightarrow F'$$

Write

$$F'_X|_H \stackrel{\text{def}}{=} H \times_{F'} F'_X; \quad (X^{\log})_H^{\text{ord}} \stackrel{\text{def}}{=} (X^{\log})_{\beta'}^{\text{ord}} \times_{F'} H$$

Recall the crystal

$$\mathcal{Q} = (Q, \nabla_Q)$$

on $\text{Crys}((X_{\beta'}^{\log})_{\mathbf{F}_p}/F')$ (from the discussion of the Hodge subspace $H \subseteq F'$ in §3.1). Let us write \mathcal{Q}_H for its restriction to $\text{Crys}((X^{\log})_H^{\text{ord}}/H)$. Thus, we have a \mathbf{P}^1 -bundle $Q_H \rightarrow X_H^{\text{ord}}$ which is equipped with a *Hodge section*

$$\sigma_1 : X_H^{\text{ord}} \rightarrow Q_H$$

On the other hand, $P \rightarrow X$ can be pulled back via γ'_X to a bundle $P_{\gamma'_X}$ on F'_X . Moreover, we have a *Hodge section*

$$\sigma_2 : F'_X \rightarrow P_{\gamma'_X}$$

Finally, by the definition of \mathcal{Q} as the *renormalized Frobenius pull-back* of $\mathcal{P}_{\gamma'}$ (cf. §3.1), we have an isomorphism

$$\Xi_{\mathcal{P}} : (P_{\gamma'_X}|_{(F'_X|_H)})_{\mathbf{Q}_p} \cong (Q_H|_{(F'_X|_H)})_{\mathbf{Q}_p}$$

(where we restrict Q_H via the projection $F'_X|_H \rightarrow X_H^{\text{ord}}$ induced by β'_X). The point here (relative to what was done in §3.1) is that by working over $F_{X'}$, which (cf. Lemma 1.9) is a sort of *universal, or tautological, Frobenius lifting* of $(X^{\log})^{\text{ord}}$, we get an isomorphism not just between crystals, but between *actual \mathbf{P}^1 -bundles* (without connection). This is important because it allows us to *compare the Hodge sections* of these \mathbf{P}^1 -bundles.

Now it follows easily from the definition of $(X^{\log})^{\text{ord}}$ (cf. the deformation argument quoted in the construction of Ψ_0^{\log} in the proof of Lemma 3.3) that there is a unique subspace

$$H'_X \subseteq F'_X|_H$$

which has the property that $\Xi_{\mathcal{P}}$ is compatible with the restrictions of σ_1 and σ_2 to H'_X . Moreover, by shrinking (i.e., replacing by an open sub-object) H'_X to some

$$H_X \subseteq F'_X|_H$$

– in such a way that the image of H_X in H still contains $\text{Im}(\xi)$, and H_X still contains the image of the section ξ_{Ψ_0} corresponding to the Frobenius lifting Ψ_0^{\log} of the proof of Lemma 3.3 – we may assume that $H_X \rightarrow X^{\text{ord}} \times_S H$ is *étale* (and surjective) and that its Kodaira-Spencer morphism (as in Definition 1.13, (3)) is an *isomorphism*. (Note that this Kodaira-Spencer morphism is an isomorphism at the image of the section ξ_{Ψ_0} precisely because Ψ_0^{\log} is “*relatively ordinary*” – as discussed in the proof of Lemma 3.3.) Thus, we see that $H_X \subseteq F'_X$ forms a *Hodge subspace* for Φ_X^{\log} , as desired.

Thus, we have proven the following result:

Theorem 3.4. *Notation as in Theorem 3.1. Then there is a unique lifting of the square of Frobenius*

$$\Phi_X^{\log} : (X^{\log})^{\text{ord}} \rightarrow (X^{\log})^{\text{ord}}$$


that covers $\Phi^{\log} : S^{\log} \rightarrow S^{\log}$ and which has the property that the induced morphism of $(\Phi_X^{\log})^* \text{Ad}(\mathcal{P}) \rightarrow \text{Ad}(\mathcal{P}) \otimes_{\mathbf{Q}_p}$ preserves the Hodge filtration. This Frobenius lifting is a very ordinary spiked Frobenius lifting of colevel $2(\chi - \Pi(1))$ (cf. Definition 1.1),

and, moreover, admits a natural choice of Hodge subspace $H_X \subseteq F'_X$ (cf. Definition 1.13) compatible with the Hodge subspace of Theorem 3.1.


In particular, one obtains on $(X^{\log})^{\text{ord}}$ (and hence, étale locally on the universal curve $\mathcal{C}_{\mathbf{Z}_p}^{\log}$ over $(\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{Z}_p}$) strong uniformizations as in Theorem 1.5, weak anabelian uniformizations as in Theorem 2.3, and a uniformizing Galois representation as in Theorem 2.11.

Pictorial Appendix


We begin with a picture of the *non-affine geometry* discussed in §2.2.



Next, we present an illustration of the (spectra of the) *total and strong Galois mantles* (cf. Definition 2.12).



Finally, we give a picture of how one might envision the *moduli* of a smooth proper curve of genus 3 parametrized by a *spiked geometry* (of *colevel* $c = 2$). Unlike the complex case, where the gluing together at each partition curve is based on giving a *twisting angle* and a *radius*, i.e., an element of $\mathbf{C}^\times = \mathbf{G}_m(\mathbf{C})$, in the p -adic case spiked case, some moduli – i.e., the *strong moduli*, of which there are $c = 2$ – are parametrized not by \mathbf{G}_m , but by the *Lubin-Tate group* \mathcal{G}_2 (cf. Theorem 1.5). Thus, it is as if some of the partition curves are not miniature copies of S^1 (i.e., the compact part of $\mathbf{G}_m(\mathbf{C})$), but rather, miniature copies of the Lubin-Tate group \mathcal{G}_2 . This is the point of view that led to the following picture:



Chapter X: Representations of the Fundamental Group of the Curve

§0. Introduction

So far, in Chapters VIII and IX, we have mainly been concerned with Galois representations of the *base*, which is some sort of *moduli stack of curves equipped with Frobenius invariant indigenous bundles*. In this Chapter, which forms the conclusion of the present work:

We “integrate” (the differential equations arising from) the various types of Frobenius invariant indigenous bundles constructed thus far and thus obtain canonical representations (Theorems 1.2 and 2.2) of the arithmetic fundamental group of the universal hyperbolic curve, which are the p -adic analogue of the canonical representation in the complex case (cf. Introduction, §0.3). These canonical representations allow one to relate the transcendental/analytic p -adic theory of the present work to the algebraic Galois action on the profinite Teichmüller group (Theorems 1.4, 2.3; cf. also §3).

As in the preceding two Chapters, we restrict ourselves to the relatively manageable *binary-ordinary* and *very ordinary spiked* cases, where we have *Galois mantles* at our disposal. The main “technology” that we use to obtain these representations is the technique of *crystalline induction*. (This is the only reason why we restrict ourselves to the binary-ordinary and very ordinary spiked cases – i.e., that crystalline induction requires the essential use of Galois mantles.) We refer to Chapter IX, §2.3, for an introduction to the technique of crystalline induction (see, especially, the Remark following Theorem 2.11). The particular crystalline induction that we carry out here is quite similar to that discussed in [Mzk1], Chapter V, §1; thus, we shall omit various technical steps that are discussed in detail in [Mzk1], Chapter V, §1.

Unfortunately, unlike in the classical ordinary case, many of the auxiliary objects (like canonical p -divisible groups and pseudo-Hecke correspondences – i.e., the theory of [Mzk1], Chapter IV) do not generalize so well to the arbitrary binary-ordinary or very ordinary spiked cases. However, in the *Lubin-Tate case*, the theory of [Mzk1], Chapter IV, generalizes immediately, so we discuss this briefly in §1.3.

§1. The Binary-Ordinary Case

Let k be a *finite field of odd characteristic* p . Let $A = W(k)$, the ring of Witt vectors with coefficients in k ; let K be its quotient field. Let S be a formally smooth, geometrically connected p -adic formal scheme over A of (constant) relative dimension d . Let us assume, moreover, that S is equipped with an étale morphism

$$S \rightarrow (\overline{\mathcal{N}}_{g,r}^{\Pi,s})_{\mathbf{Z}_p}^{\text{ord}}$$

(cf. Chapter VII, Definition 1.3) where g and r are nonnegative integers such that $2g - 2 + r \geq 1$, and Π is a *binary VF-pattern* of period ϖ . Let S^{\log} be the result of equipping S with the log structure pulled back from $\overline{\mathcal{M}}_{g,r}^{\log}$. We denote the divisor at infinity by $D \subseteq S$. Let

$$X^{\log} \rightarrow S^{\log}$$

be the pull-back of the *tautological log-curve*. Finally, let us assume that k contains \mathbf{F}_{p^ϖ} .

§1.1. The Formal \mathcal{MF}^∇ -Object

Let us denote by

$$\Phi_i^{\log} : S^{\log} \rightarrow S^{\log}$$

(where $i = 1, \dots, n$) the *canonical system of Frobenius liftings* of Chapter VII, Theorem 1.8. Write $\Lambda = \{\lambda_1, \dots, \lambda_n\}$ for the associated *multi-order* (cf. Chapter VIII, Definition 1.1). For $i = 1, \dots, n$, let

$$(\mathcal{A}[i], \nabla_{\mathcal{A}[i]}, F^\bullet(\mathcal{A}[i])) \stackrel{\text{def}}{=} (\mathcal{A}, \nabla_{\mathcal{A}}, F^\bullet(\mathcal{A}))$$

If $i \in \{1, \dots, n\}$, and $j \in \{0, \dots, \lambda_i\}$, then let

$$(\mathcal{A}[i, j], \nabla_{\mathcal{A}[i, j]}) \stackrel{\text{def}}{=} (\Phi_{S_{\mathbf{F}_p}}^*)^{j-1} \mathbf{F}^*(\mathcal{A}[i], \nabla_{\mathcal{A}[i]})$$

(if $j > 0$) and

$$(\mathcal{A}[i, j], \nabla_{\mathcal{A}[i, j]}) \stackrel{\text{def}}{=} (\mathcal{A}[i], \nabla_{\mathcal{A}[i]})$$

(if $j = 0$). Then, as we saw in Chapter IX, §2.5, Φ_i^{\log} defines a morphism

$$(\mathcal{A}[i, \lambda_i], \nabla_{\mathcal{A}[i, \lambda_i]}) \rightarrow (\mathcal{A}[i+1], \nabla_{\mathcal{A}[i+1]})$$

which becomes an *isomorphism* if one takes the p -adic completion of the PD-envelope of the left-hand side at the section defined by (cf. Chapter IX, Lemma 1.9) Φ_i^{\log} .

Write

$$D_{i,j} \stackrel{\text{def}}{=} \text{Spf}(\mathcal{A}[i, j]); \quad D_i \stackrel{\text{def}}{=} \text{Spf}(\mathcal{A}[i])$$

(where the “Spf” refers to the p -adic topology). Thus, we have a *structure morphism* $\beta_{i,j} : D_{i,j} \rightarrow S$, as well as (when $j = 0$) morphisms $\gamma_i : D_i \rightarrow S$ given by “projection to the left.” Let us equip $D_{i,j}$ with the log structure pulled back from S^{\log} . Thus, $\beta_{i,j}$ and γ_i extend to log morphisms (cf. the discussion at the beginning of Chapter IX, §1.4)

$$\beta_{i,j}^{\log} : D_{i,j}^{\log} \rightarrow S^{\log}, \quad \gamma_i^{\log} : D_i^{\log} \rightarrow S^{\log}$$

respectively. By pulling back via $\beta_{i,j}^{\log}$, we also obtain curves

$$X_{i,j}^{\log} \rightarrow D_{i,j}^{\log}$$

For $j = 0$, we denote $X_{i,j}^{\log} \rightarrow D_{i,j}^{\log}$ by $X_i^{\log} \rightarrow D_i^{\log}$.

For $i = 1, \dots, n$, let \mathcal{P}_i be the i^{th} *canonical indigenous bundle* on X^{\log} (cf. Chapter VII, Theorem 1.8). Now we execute the following procedure (cf. [Mzk1], Chapter V, the discussion preceding Theorem 1.1):

- (1) Note that $(X^{\log}, \mathcal{P}_i)$ defines a *crystal* on $\text{Crys}(X^{\log}/S^{\log})$. Thus, if we *pull-back* this crystal via $\gamma_i^{\log} : D_i^{\log} \rightarrow S^{\log}$ (i.e., from the left), then we obtain a *crystal* on

$$\text{Crys}(X^{\log}/D_i^{\log})$$

(where the structure morphism $X^{\log} \rightarrow D_i^{\log}$ is the composite of $X^{\log} \rightarrow S^{\log}$ with the diagonal embedding $S^{\log} \hookrightarrow D_i^{\log}$). We denote the \mathbf{P}^1 -bundle obtained by

evaluating this crystal on the thickening $X_i^{\log} \rightarrow D_i^{\log}$ by $Q_i \rightarrow X_i^{\log}$.

- (2) If we write $S^{\log} \times^{\text{PD}} S^{\log} \times^{\text{PD}} S^{\log}$ for the p -adic completion of the PD-envelope of the diagonal $S^{\log} \hookrightarrow S^{\log} \times_A S^{\log} \times_A S^{\log}$, then one sees easily that the pull-back of the above crystal to $\text{Crys}(X^{\log}/S^{\log} \times^{\text{PD}} S^{\log} \times^{\text{PD}} S^{\log})$ defined by the projection $S^{\log} \times^{\text{PD}} S^{\log} \times^{\text{PD}} S^{\log} \rightarrow D_i^{\log}$ to the first and second components *coincides* with the pull-back via the projection to the first and third components. Indeed, this follows from the fact that the crystal of (1) is obtained by pull-back from a crystal on $\text{Crys}(X^{\log}/S^{\log})$ via the projection $S^{\log} \times^{\text{PD}} S^{\log} \times^{\text{PD}} S^{\log} \rightarrow S^{\log}$ to the *first* component.
- (3) In particular, if we denote by $X^{\log} \times^{\text{PD}} X^{\log}$ the p -adic completion of the PD-envelope of the diagonal $X^{\log} \hookrightarrow X^{\log} \times_A X^{\log}$, then the two pull-backs of $Q_i \rightarrow X_i^{\log} = D_i^{\log} \times_{\beta_i, S^{\log}} X^{\log} \rightarrow X^{\log}$ via the projections $\pi_1, \pi_2 : X^{\log} \times^{\text{PD}} X^{\log} \rightarrow X^{\log}$ *coincide* – indeed, this follows by applying the observation of (2) above to the commutative diagram (for $\alpha = 1, 2$)

$$\begin{array}{ccc}
 S^{\log} \times^{\text{PD}} X^{\log} \times^{\text{PD}} X^{\log} = D_i^{\log} \times_{\beta_i, S^{\log}, \pi_{\alpha}} (X^{\log} \times^{\text{PD}} X^{\log}) & \longrightarrow & S^{\log} \times^{\text{PD}} S^{\log} \times^{\text{PD}} S^{\log} \\
 \downarrow 1 \times \pi_{\alpha} & & \downarrow \pi_{1, \alpha+1} \\
 X_i^{\log} = D_i^{\log} \times_{\beta_i, S^{\log}} X^{\log} & \longrightarrow & D_i^{\log}
 \end{array}$$

(where the vertical morphism on the right is the projection to the factors numbered 1 and $\alpha + 1$), which may be regarded as defining a morphism from a *thickening* of

$X^{\log} \rightarrow S^{\log} \times^{\text{PD}} S^{\log} \times^{\text{PD}} S^{\log}$ to a thickening of $X^{\log} \rightarrow D_i^{\log}$, i.e., a morphism from an object of $\text{Crys}(X^{\log}/S^{\log} \times^{\text{PD}} S^{\log} \times^{\text{PD}} S^{\log})$ to an object of $\text{Crys}(X^{\log}/D_i^{\log})$ that covers the projection $\pi_{1, \alpha+1} : S^{\log} \times^{\text{PD}} S^{\log} \times^{\text{PD}} S^{\log} \rightarrow D_i^{\log}$.

- (4) Thus, by *composing* $Q_i \rightarrow X_i^{\log}$ with the projection $X_i^{\log} \rightarrow X^{\log}$, we obtain a relative (p -adic formal log) scheme $Q_i \rightarrow X^{\log}$, which is equipped with a natural *integrable* (by an argument similar to that of (3) above) *connection* ∇_{Q_i} relative to $X^{\log} \rightarrow \text{Spec}(A)$. Moreover,

$$(Q_i, \nabla_{Q_i})$$

has the structure of an object – more precisely, a \mathbf{P}^1 -bundle – over $(D_i, \nabla_{D_i}) \times_S X$. Note that this \mathbf{P}^1 -bundle is no longer indigenous, but becomes indigenous when restricted to the diagonal $\Delta_i \subseteq D_i$. Thus, (Q_i, ∇_{Q_i}) is equipped with a *Hodge subspace* given by taking the image of the Hodge section $\Delta_i \rightarrow Q_i|_{\Delta_i}$ defined on the diagonal $\Delta_i \subseteq D_i$.

The procedure carried out above is entirely similar to that carried out in the discussion preceding [Mzk1], Chapter V, Theorem 1.1; we refer to *loc. cit.* for more details.

Note that since $D_{i,1}$ is the *renormalized Frobenius pull-back* of D_i (cf. Chapter IX, Definition 1.10), it follows that the morphism $(D_{i,1})_{\mathbf{F}_p} \rightarrow (\Phi_{S_{\mathbf{F}_p}}^* D_i)_{\mathbf{F}_p}$ factors through the diagonal

$$(\Phi_{S_{\mathbf{F}_p}}^* \Delta_i)_{\mathbf{F}_p} \subseteq (\Phi_{S_{\mathbf{F}_p}}^* D_i)_{\mathbf{F}_p}$$

Thus,

$$(\Phi_{X_{\mathbf{F}_p}}^* (Q_i, \nabla_{Q_i}) \times_{(\Phi_{S_{\mathbf{F}_p}}^* D_i)} D_{i,1})_{\mathbf{F}_p}$$

is equipped with a *horizontal Hodge section*. In particular, we may form the *renormalized Frobenius pull-back* of $(Q_{i,0}, \nabla_{Q_{i,0}}) \stackrel{\text{def}}{=} (Q_i, \nabla_{Q_i})$. If we do this, and then apply various naive Frobenius pull-backs, then we obtain objects (for $j = 1, \dots, \lambda_i$)

$$\begin{aligned} (Q_{i,j}, \nabla_{Q_{i,j}}) &\stackrel{\text{def}}{=} (\Phi_{X_{\mathbf{F}_p}}^*)^{j-1} \mathbf{F}^* (\Phi_{S_{\mathbf{F}_p}}^* (Q_i, \nabla_{Q_i}) \times_{(\Phi_{S_{\mathbf{F}_p}}^* D_i)} D_{i,1}) \\ &= (\Phi_{X_{\mathbf{F}_p}/S_{\mathbf{F}_p}}^*)^{j-1} \mathbf{F}^* (\Phi_{S_{\mathbf{F}_p}}^* (Q_i, \nabla_{Q_i}) \times_{(\Phi_{S_{\mathbf{F}_p}}^* D_i)} D_{i,j}) \end{aligned}$$

on X^{\log} , equipped with an auxiliary structure of \mathbf{P}^1 -bundle over

$$(D_{i,j}, \nabla_{D_{i,j}}) \times_S X$$

Moreover, by the “compatibility condition” of Chapter VII, Theorem 1.8, we have *isomorphisms*

$$(Q_{i+1}, \nabla_{Q_{i+1}}) \cong (Q_{i,\lambda_i}, \nabla_{Q_{i,\lambda_i}})$$

which constitute a *Frobenius action* on the $(Q_{i,j}, \nabla_{Q_{i,j}})$ (for $i = 1, \dots, n$; $j = 0, \dots, \lambda_i - 1$). In summary, we have the following result (cf. [Mzk1], Chapter V, Theorem 1.1):

Theorem 1.1. *The object*

$$\prod_{i,j} \mathcal{Q}_{i,j}$$

(where the product is fibered over X^{\log} and runs over $i = 1, \dots, n$; $j = 0, \dots, \lambda_i - 1$) is equipped with all the structures necessary for a “formal” \mathcal{MF}^∇ -object on X^{\log} : full logarithmic connections (relative to $X^{\log} \rightarrow \mathrm{Spec}(A)$), Hodge subspaces, and a Frobenius action. Moreover, it is also equipped with a structure of \mathbf{P}^1 -bundle over the formal \mathcal{MF}^∇ -object

$$\left(\prod_{i,j} D_{i,j}^{\log} \right) \times_{S^{\log}} X^{\log}$$

(where the indexed product is fibered over S^{\log} and runs over $i = 1, \dots, n$; $j = 0, \dots, \lambda_i - 1$, and the implicit morphism from this indexed product to S^{\log} is the structure morphism) on X^{\log} . Finally, this structure of \mathbf{P}^1 -bundle is compatible with the various full logarithmic connections, Hodge subspaces, and Frobenius actions involved.

§1.2. The Crystalline Induced Representation

Now, just as in Chapter IX, §2.3 (of the present work) and [Mzk1], Chapter V, Theorem 1.4, we use the technique of base-change to the *Galois mantle* (cf. Chapter IX, Definition 2.13) to pass from the formal \mathcal{MF}^∇ -object of Theorem 1.1 to a Galois representation. Since the technique of proof is entirely similar to that of Chapter IX, §2.3, and [Mzk1], Chapter V, §1, we shall omit various details.

First, let

$$Z_K^{\log} \rightarrow S_K^{\log}$$

be the finite log étale covering of the discussion preceding Chapter IX, Definition 2.13. Recall the *morphism* (of the discussion following Chapter IX, Definition 2.13):

$$(\mathcal{A}[i, j])(B^+(S^{\log})) \rightarrow (\mathcal{A}_\Phi^{\mathrm{Gal}})_{B^+}^{F^{N_{i,j}}}$$

where

$$(\mathcal{A}_\Phi^{\mathrm{Gal}})_{B^+} \stackrel{\mathrm{def}}{=} B^+(S^{\log}) \otimes_{\mathcal{O}_\varpi} \mathcal{A}_\Phi^{\mathrm{Gal}}$$

and $N_{i,j} = j + \sum_{a=1}^{i-1} \lambda_a$. This morphism *respects Hodge filtrations, Frobenius, and the action of $\Pi_{Z^{\log}}$* .

Let

$$U^{\log} \subseteq X^{\log}$$

be a “small” (p -adic) affine open (cf. Chapter IX, §2.3). If we *base-change* by means of the above morphism the \mathbf{P}^1 -bundle obtained by evaluating (on the thickening defined by $B^+(U^{\log})$) the crystal in \mathbf{P}^1 -bundles given by:

$$(Q_{i,j}, \nabla_{Q_{i,j}}) \rightarrow (D_{i,j}^{\log} \stackrel{\text{def}}{=} \text{Spf}(\mathcal{A}[i,j]), \nabla_{D_{i,j}^{\log}}) \times_{S^{\log}} X^{\log}$$

(where the morphisms to S^{\log} in the product are the structure morphisms) then we obtain a \mathbf{P}^1 -bundle

$$Q_{i,j}^{\text{Gal}}|_U \rightarrow \text{Spf}\{(\mathcal{A}_{\Phi}^{\text{Gal}})_{B^+}^{F^{N_{i,j}}} \otimes_{B^+(S^{\log})} B^+(U^{\log})\}$$

Moreover, if we base-change *this* \mathbf{P}^1 -bundle by

$$B^+(U^{\log}) \rightarrow \mathcal{O}_{\tilde{U}_{\mathbf{F}_p}}$$

(where we use analogous notation to the notation of Chapter IX, §2.3 – i.e., \tilde{U} is the “ U^{\log} -version” of \tilde{S}), and then further base-change by the *Frobenius morphism* of $\mathcal{O}_{\tilde{U}_{\mathbf{F}_p}}$, then the resulting \mathbf{P}^1 -bundle

$$(Q_{i,j}^{\text{Gal}}|_U \otimes_{B^+(U^{\log})} \mathcal{O}_{\tilde{U}_{\mathbf{F}_p}})^F$$

admits a Hodge section if $j = 0$ (respectively, *has an empty Hodge subspace* if $j > 0$). Indeed, that the Hodge subspace is empty if $j > 0$ follows immediately from the definitions. The existence of the Hodge section if $j = 0$ follows from the discussion preceding Theorem 1.1 in §1.1 (and the fact that $(\mathcal{A}_{\Phi}^{\text{Gal}})^{F^{N_{i,j}}}$ consists of *Frobenius invariant* elements of the result of evaluating $\mathcal{A}[i,j]$ on some thickening).

Thus, it follows that we may form the *renormalized (respectively, naive) Frobenius pull-back* of the \mathbf{P}^1 -bundle $Q_{i,j}^{\text{Gal}}|_U$ if $j = 0$ (respectively, $j > 0$): i.e., we have

$$\mathbf{F}^*(Q_{i,0}^{\text{Gal}}|_U) \cong Q_{i,1}^{\text{Gal}} \quad (\text{respectively, } \Phi^*(Q_{i,j}^{\text{Gal}}|_U) \cong Q_{i,j+1}^{\text{Gal}})$$

(cf. the discussion preceding Chapter IX, Lemma 2.9). The point here is that (cf. Chapter IX, §2.3; [Mzk1], Chapter V, §1):

After the base-change carried out above, the \mathbf{P}^1 -bundle that we obtain is essentially the projectivization of a rank two vector bundle with a one-step (respectively, zero-step) Hodge filtration if $j = 0$ (respectively, $j > 0$).

That is to say, the number of steps in the Hodge filtration is *finite* and $\leq p-2$. Thus, just as in [Mzk1], Chapter V, Proposition 1.3, we obtain that (at least after further base-change via $\mathcal{A}_{\Phi}^{\text{Gal}} \rightarrow \widehat{\mathcal{A}}_{\Phi}^{\text{Gal}}$ – where $\widehat{\mathcal{A}}_{\Phi}^{\text{Gal}}$ is the *completion* of $\mathcal{A}_{\Phi}^{\text{Gal}}$ with respect to the filtration defined by the divided powers of the augmentation ideal) the theory of [Falt1], §2 (cf. especially the proof of [Falt1], §2, Theorem 2.4) already implies that:

The “ $F^0(-)^{F=1}$ ” of the product of the \mathbf{P}^1 -bundles

$$Q_{i,j}^{\text{Gal}}|_U \rightarrow \text{Spf}\{(\mathcal{A}_{\Phi}^{\text{Gal}})_{B^+}^{F^{N_{i,j}}} \otimes_{B^+(S^{\log})} B^+(U^{\log})\}$$

(where $i = 1, \dots, n$; $j = 0, \dots, \lambda_i - 1$) is “of the expected size.”

Here, the implicit “descent” from $\widehat{\mathcal{A}}_{\Phi}^{\text{Gal}}$ back down to $\mathcal{A}_{\Phi}^{\text{Gal}}$ is entirely analogous to Chapter IX, Lemma 2.9 (i.e., one takes the inclusion $\widehat{\mathcal{A}}^{\text{st,Gal}} \hookrightarrow \mathcal{R}$ of *loc. cit.* to be the identity $\widehat{\mathcal{A}}^{\text{Gal}} = \widehat{\mathcal{A}}^{\text{Gal}}$).

Thus, by gluing together the Frobenius invariants obtained from various opens U^{\log} , we obtain the following result (cf. [Mzk1], Chapter V, Theorem 1.4):

Theorem 1.2. *Let*

$$\mathcal{A}_{\Phi}^{\text{Gal}}$$

be the Galois mantle (i.e., the $\mathcal{O}_{\omega}[\Pi_{Z^{\log}}]$ -algebra of Chapter IX, Definition 2.13) associated to the canonical binary-ordinary Frobenius system under consideration (cf. Chapter VIII, Theorem 3.1). Let

$$\mathbf{A}_{\Phi}^{\text{Gal}} \stackrel{\text{def}}{=} \text{Spf}(\mathcal{A}_{\Phi}^{\text{Gal}})$$

(where “Spf” is taken with respect to the p -adic topology). Then associated to the formal \mathcal{MF}^{∇} -object of Theorem 1.1, we have a \mathbf{P}^1 -bundle

$$\mathbf{P}_{\Phi}^{\text{Gal}} \rightarrow \mathbf{A}_{\Phi}^{\text{Gal}}$$

equipped with a continuous action by $\pi_1((X_Z^{\log})_K)$ (where $X_Z^{\log} \stackrel{\text{def}}{=} X^{\log} \times_{S^{\log}} Z^{\log}$) which is compatible with the action of $\Pi_{Z^{\log}} \stackrel{\text{def}}{=} \pi_1(Z_K^{\log})$ on $\mathbf{A}_{\Phi}^{\text{Gal}}$.

Remark. Once one has this Galois \mathbf{P}^1 -bundle $\mathbf{P}_{\Phi}^{\text{Gal}}$, one can carry out a discussion of the relationship between it and the canonical group-theoretic or affine parameters of Chapter VIII in a fashion entirely analogous to that of the discussion preceding [Mzk1], Chapter V, Theorem 1.7. (In particular, one obtains a binary-ordinary analogue of the observation (*) immediately following [Mzk1], Chapter V, Theorem 1.7.) Since the proofs and results are entirely similar to those of [Mzk1], Chapter V, §1, we leave it to the reader to make them explicit.

§1.3. The Lubin-Tate Case

Unfortunately, unlike the classical ordinary case, in general (even under the assumption of binary-ordinariness), it is unpleasant to try to work out a theory of canonical p -divisible groups and pseudo-Hecke correspondences on X^{\log} (as in [Mzk1], Chapter IV). However, at least in the Lubin-Tate case, the results of [Mzk1], Chapter IV, generalize quite directly, so we will pause here to take a brief look at what happens. Thus, in the rest of this subsection, we assume that Π is of pure tone $\varpi \geq 2$. In particular, we may assume that the log structures on X^{\log} and S^{\log} are trivial.

Let $\alpha \in S(A)$ be an A -valued point. Let

$$X_{\alpha} \rightarrow \text{Spec}(A)$$

be the restriction of $X \rightarrow S$ to α . Then we have the following result:

Theorem 1.3. *The following are equivalent:*

- (1) *the point α is canonical (cf. Chapter VIII, Definition 1.3 – in the case where “ σ ” is the identity permutation of the set of one element) for the canonical Frobenius lifting Φ^{\log} ;*
- (2) *X_{α} admits an indigenous bundle (P, ∇_P) such that*

$$((\Phi_{(X_{\alpha})_{\mathbf{F}_p}})^*)^{\varpi-1} \mathbf{F}^*(P, \nabla_P) \cong (P, \nabla_P)$$

Moreover, if α is canonical, then X_{α} admits a canonical dual crystalline (cf. [Falt1], §2) representation

$$\rho : \pi_1((X_\alpha)_K) \rightarrow \mathrm{GL}_2^\pm(\mathcal{O}_\varpi) \stackrel{\mathrm{def}}{=} (\mathrm{GL}_2/\{\pm 1\})(\mathcal{O}_\varpi)$$

as well as a canonical p -divisible group G (up to ± 1) equipped with an \mathcal{O}_ϖ -action whose associated Tate module gives rise to the representation ρ .

Suppose that X_α is canonical. Let

$$Y_\alpha \rightarrow X_\alpha$$

be the scheme parametrizing isogenies $G \rightarrow H$ whose kernel is a cyclic \mathcal{O}_ϖ -module of order $q \stackrel{\mathrm{def}}{=} p^\varpi$. Then $(Y_\alpha)_{\mathbf{F}_p}$ admits the same sort of analysis (cf. [Mzk1], Chapter IV, Theorem 2.5) as in the classical ordinary case: that is, $(Y_\alpha)_{\mathbf{F}_p}$ is isomorphic to the “union of the graph of $\Phi_{(X_\alpha)_{\mathbf{F}_p}}^\varpi$ and its transpose, joined at the supersingular points.” Moreover, the supersingular locus of the canonical indigenous bundle on $(X_\alpha)_{\mathbf{F}_p}$ is étale over k (cf. [Mzk1], Chapter IV, Proposition 3.2) and admits a natural $\mathbf{F}_{p^{2\varpi}}$ -rational structure.

Although we shall not write out the details, X_α also admits a canonical “pseudo-Hecke correspondence” (cf. [Mzk1], Chapter IV, Definition 3.3) that lifts “the graph of $\Phi_{(X_\alpha)_{\mathbf{F}_p}}^\varpi$ and its transpose, joined at the supersingular points.” Finally, just as in the classical ordinary case, one has a *geometric criterion for canonicity* (whose statement and proof are entirely analogous to the classical ordinary case), just as in [Mzk1], Chapter IV, Theorem 4.17.

The reader is also advised to refer to the theory of [Ih1-4] for a thorough discussion of the *Shimura curve case*.

§1.4. Relation to the Profinite Teichmüller Group

In this subsection, let us write

$$\mathcal{C} \rightarrow \mathcal{M} \stackrel{\mathrm{def}}{=} (\overline{\mathcal{M}}_{g,r}^{\mathrm{log}})_{\mathbf{Z}_p}$$

for the universal log-curve over the moduli stack $(\overline{\mathcal{M}}_{g,r}^{\mathrm{log}})_{\mathbf{Z}_p}$ of r -pointed stable log-curves of genus g over \mathbf{Z}_p . We shall use subscripted “ \mathbf{Q}_p ’s” to denote the result of base-changing objects over \mathbf{Z}_p via $\mathbf{Z}_p \rightarrow \mathbf{Q}_p$. Let us denote by η the generic point of \mathcal{M} , and by $\bar{\eta}$ the spectrum of a field $k(\bar{\eta})$ such that $k(\bar{\eta})$ is an algebraic closure of $k(\eta)$. Let us write

$$\mathcal{C}_{\bar{\eta}} \stackrel{\mathrm{def}}{=} \mathcal{C} \times_{\mathcal{M}} \bar{\eta}$$

for the geometric fiber of $\mathcal{C} \rightarrow \mathcal{M}$.

In the following, we would like to consider the *fundamental group* of various objects. Although, strictly speaking, in order to specify the fundamental group of a scheme/stack, one must specify a base-point of the scheme/stack, we will omit the mention of base-points in the following discussion (since the issues that arise from the ambiguity of the choice of base-point will be irrelevant). Also, in the following, we will sometimes consider “ π_1 ” of schemes/stacks which are *not connected*. The reader should interpret this sort of π_1 in a “topos-theoretic sense” as follows: Since the only sense in which we will need such π_1 ’s is for the categories of profinite sets with π_1 -action that they define, the reader should understand the notion of “a set with the action of π_1 of a non-connected scheme” to mean “a collection of sets, one for each connected component of the scheme, equipped with an action of π_1 of that connected component.”

Thus, with these conventions concerning π_1 , we obtain an *exact sequence of profinite groups*

$$1 \rightarrow \pi_1(\mathcal{C}_{\overline{\eta}}) \rightarrow \pi_1(\mathcal{C}_{\mathbf{Q}_p}) \rightarrow \pi_1(\mathcal{M}_{\mathbf{Q}_p}) \rightarrow 1$$

(cf. [Oda], [Naka]). Here,

$$\pi_1(\mathcal{C}_{\overline{\eta}})$$

is the *profinite completion of the (topological) fundamental group of a Riemann surface of genus g with r punctures*. Moreover, $\pi_1(\mathcal{M}_{\mathbf{Q}_p})$ fits into an exact sequence

$$1 \rightarrow \pi_1(\mathcal{M}_{\overline{\mathbf{Q}_p}}) \rightarrow \pi_1(\mathcal{M}_{\mathbf{Q}_p}) \rightarrow \text{Gal}(\overline{\mathbf{Q}_p}/\mathbf{Q}_p) \rightarrow 1$$

and $\pi_1(\mathcal{M}_{\overline{\mathbf{Q}_p}})$ is the *profinite completion of the Teichmüller group for Riemann surfaces of genus g with r punctures*. Of course, all of this (i.e., the above two exact sequences) could be done over \mathbf{Q} instead of \mathbf{Q}_p , but we work over \mathbf{Q}_p since (in this subsection):

Our aim is to relate these exact sequences of algebraic fundamental groups to the (highly analytic) p-adic theory developed so far.

Observe that the first exact sequence above induces an *outer action* of $\pi_1(\mathcal{M}_{\mathbf{Q}_p})$ on $\pi_1(\mathcal{C}_{\overline{\eta}})$, i.e., a homomorphism

$$\pi_1(\mathcal{M}_{\mathbf{Q}_p}) \rightarrow \text{Out}(\pi_1(\mathcal{C}_{\overline{\eta}})) \stackrel{\text{def}}{=} \text{Aut}(\pi_1(\mathcal{C}_{\overline{\eta}}))/\pi_1(\mathcal{C}_{\overline{\eta}})$$

(where the quotient on the right is the quotient by the “inner automorphisms” of $\pi_1(\mathcal{C}_{\overline{\eta}})$).

Now recall the ring $\mathcal{O}_{\varpi} \stackrel{\text{def}}{=} W(\mathbf{F}_{p^\varpi})$. Let us write

$$\mathbf{Rep} \stackrel{\text{def}}{=} \text{Rep}(\pi_1(\mathcal{C}_{\overline{\eta}}), PGL_2(\mathcal{O}_{\varpi}))$$

for the (profinite) set of *continuous homomorphisms* $\pi_1(\mathcal{C}_{\overline{\eta}}) \rightarrow PGL_2(\mathcal{O}_{\varpi})$ considered up to composition with an inner automorphism of $PGL_2(\mathcal{O}_{\varpi})$. Note that this set is equipped with a natural topology. Moreover, the outer action of $\pi_1(\mathcal{M}_{\mathbf{Q}_p})$ on $\pi_1(\mathcal{C}_{\overline{\eta}})$ induces a *continuous action* of $\pi_1(\mathcal{M}_{\mathbf{Q}_p})$ on \mathbf{Rep} . Then it follows from the definition of the fundamental group that \mathbf{Rep} corresponds to some *infinite covering* (i.e., inverse limit of finite log étale coverings)

$$\mathcal{R}_{\mathbf{Q}_p} \rightarrow \mathcal{M}_{\mathbf{Q}_p}$$

Let us denote the *normalization* of \mathcal{M} in $\mathcal{R}_{\mathbf{Q}_p}$ by \mathcal{R} . Thus, $\mathcal{R}_{\mathbf{Q}_p} = \mathcal{R} \otimes_{\mathbf{Z}_p} \mathbf{Q}_p$.

Next, let us write

$$\mathcal{N} \stackrel{\text{def}}{=} (\overline{\mathcal{N}}_{g,r}^{\Pi,s})_{\mathbf{Z}_p}^{\text{ord},\log}$$

Recall that we have a natural morphism $\mathcal{N} \rightarrow \mathcal{M}$. Next, recall the “set-theoretic Galois representation” \mathcal{U} of Chapter VIII, Definition 1.4. This set \mathcal{U} is a profinite set with a continuous $\pi_1(\mathcal{N}_{\mathbf{Q}_p})$ -action. Thus, it corresponds to some *infinite covering* (i.e., inverse limit of finite log étale coverings)

$$\mathcal{N}[\infty]_{\mathbf{Q}_p} \rightarrow \mathcal{N}_{\mathbf{Q}_p}$$

We will denote the normalization of \mathcal{N} in $\mathcal{N}[\infty]_{\mathbf{Q}_p}$ by

$$\mathcal{N}[\infty] \rightarrow \mathcal{N}$$

It follows immediately from the definitions that $Z_{\mathbf{Q}_p}^{\log}$ (where Z^{\log} is as in Theorem 1.2 in the case where we take “ S^{\log} ” to be \mathcal{N} itself) is an *intermediate covering* between $\mathcal{N}_{\mathbf{Q}_p}$ and $\mathcal{N}[\infty]_{\mathbf{Q}_p}$. Let us write \tilde{Z}^{\log} for the *normalization* of \mathcal{N} in $Z_{\mathbf{Q}_p}^{\log}$. Thus, we have morphisms

$$\pi_1(\mathcal{N}[\infty]_{\mathbf{Q}_p}) \rightarrow \pi_1(Z_{\mathbf{Q}_p}^{\log}) \rightarrow \pi_1(\mathcal{N}_{\mathbf{Q}_p})$$

In particular, we have a natural action of $\pi_1(\mathcal{N}[\infty]_{\mathbf{Q}_p})$ on $\mathbf{A}_{\Phi}^{\text{Gal}}$ (cf. Theorem 1.2). Moreover, if one interprets Chapter VIII, Theorem 2.12,

in the present context, one sees that the content of that theorem is essentially that $\mathbf{A}_\Phi^{\text{Gal}}$ admits a *tautological* \mathcal{O}_ϖ -valued point

$$\alpha_\infty \in \mathbf{A}_\Phi^{\text{Gal}}(\mathcal{O}_\varpi)$$

which is fixed by $\pi_1(\mathcal{N}[\infty]_{\mathbf{Q}_p})$. Thus, if one restricts the “Galois \mathbf{P}^1 -bundle” of Theorem 1.2 to the point α_∞ (and uses the fact that the automorphisms of \mathbf{P}^1 are given schematically by “ PGL_2 ”), we obtain from the definition of \mathbf{Rep} that we have a natural *classifying morphism*

$$\kappa_{\mathbf{Q}_p} : \mathcal{N}[\infty]_{\mathbf{Q}_p} \rightarrow \mathcal{R}_{\mathbf{Q}_p}$$

for the restriction to α_∞ of the “Galois \mathbf{P}^1 -bundle” of Theorem 1.2. Moreover, it follows from the definition (as normalizations) of $\mathcal{N}[\infty]$ and \mathcal{R} that $\kappa_{\mathbf{Q}_p}$ extends to a morphism

$$\kappa : \mathcal{N}[\infty] \rightarrow \mathcal{R}$$

Note that since \mathcal{N} is (as an object) a p -adic formal (log) stack over \mathbf{Z}_p , whereas \mathcal{M} is simply an algebraic (log) stack over \mathbf{Z}_p , it will be useful in comparing $\mathcal{N}[\infty]$ and \mathcal{R} to work with the p -adic completion $\widehat{\mathcal{M}}$ of \mathcal{M} . Since \mathcal{M} is proper over \mathbf{Z}_p , it follows that finite log étale coverings of $\widehat{\mathcal{M}}_{\mathbf{Q}_p}$ are equivalent to finite log étale coverings of $\mathcal{M}_{\mathbf{Q}_p}$. Let us write $\widehat{\mathcal{R}} \stackrel{\text{def}}{=} \mathcal{R} \times_{\mathcal{M}} \widehat{\mathcal{M}}$. (Note: $\widehat{\mathcal{R}}$ is a limit of p -adically complete finite coverings of $\widehat{\mathcal{M}}$, not a p -adic completion of a limit of finite coverings. In particular, (in general) $\widehat{\mathcal{R}}$ is not p -adically complete.) Thus, it follows from the definitions that κ defines morphisms

$$\widehat{\kappa} : \mathcal{N}[\infty] \rightarrow \widehat{\mathcal{R}}; \quad \widehat{\kappa}_{\mathbf{Q}_p} : \mathcal{N}[\infty]_{\mathbf{Q}_p} \rightarrow \widehat{\mathcal{R}}_{\mathbf{Q}_p}$$

The purpose of the present subsection is to prove the following result:

Theorem 1.4. *We have a natural commutative diagram:*

$$\begin{array}{ccc} \mathcal{N}[\infty] & \xrightarrow{\widehat{\kappa}} & \widehat{\mathcal{R}} \\ \downarrow & & \downarrow \\ \mathcal{N} & \longrightarrow & \widehat{\mathcal{M}} \end{array}$$

Moreover, the morphism $\widehat{\kappa}$ – which is induced by the classifying morphism for the restriction to the tautological point $\alpha_\infty \in \mathbf{A}_\Phi^{\text{Gal}}(\mathcal{O}_\varpi)$ of the “Galois \mathbf{P}^1 -bundle” of Theorem 1.2 – is an open immersion.

Proof. Thus, it remains to prove that $\widehat{\kappa}$ is an *open immersion*. Before beginning the proof, let us make some basic observations. First of all, since $\mathcal{N}[\infty]$ and $\widehat{\mathcal{R}}$ are both \mathbf{Z}_p -flat and (log) *étale* over $\widehat{\mathcal{M}}$ in characteristic zero, it is clear that to show that $\widehat{\kappa}$ is an *open immersion*, it suffices to show that it is an *immersion*. To prove that $\widehat{\kappa}$ is an immersion essentially means that distinct “points” (in some sense) of $\mathcal{N}[\infty]$ that map to the same “point” of $\widehat{\mathcal{M}}$ are mapped by $\widehat{\kappa}$ to distinct “points” of $\widehat{\mathcal{R}}$. Moreover, distinct “points” of $\mathcal{N}[\infty]$ that map to the same “point” of $\widehat{\mathcal{M}}$ have, so to speak, “three levels of distinction”:

- (1) They map to *distinct points* in \mathcal{N} .
- (2) They map to the same point in \mathcal{N} , but *distinct points* in the intermediate covering \widetilde{Z}^{\log} (between \mathcal{N} and $\mathcal{N}[\infty]$).
- (3) They map to the same point of \widetilde{Z}^{\log} , but are still *distinct as points* of $\mathcal{N}[\infty]$.

The proof of Theorem 1.4 will proceed (roughly speaking) by showing successively that distinct points of $\mathcal{N}[\infty]$ that map to the same point of $\widehat{\mathcal{R}}$ by $\widehat{\kappa}$ are, in fact, identical relative to the above three levels of distinction.

Let us begin with the *first level of distinction* (i.e., (1) above). Thus, we take two points

$$\xi_1, \xi_2 \in \mathcal{N}[\infty](\overline{\mathbf{F}}_p)$$

satisfying $\widehat{\kappa}(\xi_1) = \widehat{\kappa}(\xi_2)$ and show that ξ_1 and ξ_2 have the same image in $\mathcal{N}(\overline{\mathbf{F}}_p)$. First, let us observe that although \mathcal{N} is (in general) *not finite* over $\widehat{\mathcal{M}}$, after *étale localization* on $\widehat{\mathcal{M}}$ (via, say, some *étale* morphism $S \rightarrow \widehat{\mathcal{M}}$), we may assume that each fiber of $\mathcal{N}_{\mathbf{F}_p} \rightarrow \widehat{\mathcal{M}}_{\mathbf{F}_p}$ is contained in an open and closed sub-object of $\mathcal{N}|_S$ which is a *disjoint union of a finite number of copies of S* . Thus, by *étale localizing* in this fashion, working over $B^+(S)$ (cf. §1.2), and using the fact that since $\mathcal{N}[\infty]$ and $\widehat{\mathcal{R}}$ are defined as *normalizations*, $\widehat{\kappa}(\xi_1) = \widehat{\kappa}(\xi_2)$ implies that the entire “sheets” corresponding to ξ_1 and ξ_2 map *generically* to the same “sheet” of $\widehat{\mathcal{R}}$, we thus see that we are reduced to the following issue:

- (*) We have two indigenous bundles on $\mathcal{C}|_{\mathrm{Spf}(B^+(S))}$ which are Π -invariant. These indigenous bundles (together with their associated “auxiliary bundles” – cf. Chapter VI, Definition 1.19) both define \mathcal{MF}^∇ -objects, hence (cf. §1.2; [Falt1], §2) representations of $\pi_1(\mathcal{C}_{\overline{\eta}})$. Then under the assumption that these two representations

coincide, we would like to conclude that the two indigenous bundles, as well as the various associated “auxiliary” bundles *coincide modulo p* (since this is precisely the data that defines a point of $\mathcal{N}_{\mathbb{F}_p}$, which was formed from the shifted VF-stacks of Chapter III, §1.3).

Note that, although in [Falt1], Theorem 2.6, it is stated that the functor “ \mathbf{D} ” (which converts \mathcal{MF}^∇ -objects into Galois representations) is *fully faithful*, the result that we need here is somewhat different since here we only know that the representations of the “geometric fundamental group” $\pi(\mathcal{C}_{\overline{\eta}})$ coincide. In fact, what we need here is essentially the “*relative comparison theorem*” (i.e., [Falt1], Theorem 6.2) in the case where (in the notation of [Falt1], Theorem 6.2) “ $b = 0$.” That is to say, this result implies, in particular, that the existence of an isomorphism between the corresponding representations of the *geometric* fundamental group implies the existence of an isomorphism between the corresponding \mathbf{P}^1 -bundles with *relative* (i.e., for $\mathcal{C}|_{\mathrm{Spf}(B^+(S))} \rightarrow \mathrm{Spf}(B^+(S))$) connection. Sorting through the definitions shows that this completes the verification of $(*)_1$, and thus shows that $\widehat{\kappa}$ *distinguishes fibers over \mathcal{N}* .

Remark. Some mathematicians have called into question the validity of the techniques used to prove [Falt1], Theorem 6.2. In the present case, however, where “ $b = 0$,” it is possible to check the result fairly directly (using Faltings’ theory of almost étale extensions). We leave this as an exercise to the reader.

Now that we have shown that $\widehat{\kappa}$ distinguishes fibers over \mathcal{N} , it follows that by étale localization on $\widehat{\mathcal{M}}$ and \mathcal{N} , we are free to act as though $\widehat{\kappa}$ is a morphism between inverse limits of *finite* coverings over some base S . (Here, S will lie over \mathcal{N} in such a way that the morphism of (non-logarithmic!) algebraic stacks underlying $S \rightarrow \mathcal{M}$ is étale.) Thus, we obtain an S -morphism

$$\widehat{\kappa}_S : \mathcal{N}[\infty]_S \rightarrow \widehat{\mathcal{R}}_S$$

(where $\mathcal{N}[\infty]_S$ (respectively, $\widehat{\mathcal{R}}_S$) is the object obtained from $\mathcal{N}[\infty]$ (respectively, $\widehat{\mathcal{R}}$) by the above-mentioned étale localization on \mathcal{N} (respectively, $\widehat{\mathcal{M}}$)) of inverse limits of finite coverings of S , and we wish to show that $\widehat{\kappa}_S$ is an *immersion*. Moreover, since $\mathcal{N}[\infty]$ and \mathcal{R} are both defined as *normalizations*, it suffices to show that

$$\widehat{\kappa}_{S_{\mathbf{Q}_p}} \stackrel{\text{def}}{=} \widehat{\kappa}_S \otimes \mathbf{Q}_p : \mathcal{N}[\infty]_{S_{\mathbf{Q}_p}} \rightarrow \widehat{\mathcal{R}}_{S_{\mathbf{Q}_p}}$$

is an *immersion* (of inverse limits of finite log étale coverings of $S_{\mathbf{Q}_p}$). In another words, we have reduced the problem to a problem which may be phrased *entirely in terms of (profinite) $\pi_1(S_{\mathbf{Q}_p})$ -sets*.

Next, observe that the covering $\mathcal{N}[\infty]_S \rightarrow S$ factors as

$$\mathcal{N}[\infty]_S \rightarrow \tilde{Z}_S^{\log} \rightarrow S$$

Thus, our next step is to show that $\widehat{\kappa}_{S_{\mathbf{Q}_p}}$ distinguishes fibers over $\tilde{Z}_{S_{\mathbf{Q}_p}}^{\log} = Z_{S_{\mathbf{Q}_p}}^{\log}$. On the other hand, the covering $Z_{S_{\mathbf{Q}_p}}^{\log} \rightarrow S_{\mathbf{Q}_p}$ is, by definition (cf. the discussion of Chapter IX, §2.5), the covering parametrizing splittings modulo p of a certain Galois torsor. Moreover, this Galois torsor is precisely the Galois object corresponding to the “geometric portion” of \mathcal{A} (i.e., the sub-object of \mathcal{A} which is isomorphic to the object \mathcal{P} of Chapter VIII, §2.2). Thus, what we must show here is that:

(\ast_2) Infinitesimal deformations modulo p of splittings of the Galois torsor of Chapter IX, §2.5, correspond to (i.e., induce) *nonzero* deformations of the $PGL_2(\mathcal{O}_{\varpi})$ -valued representation of $\pi_1(\mathcal{C}_{\overline{\eta}})$ arising from (restricting to α_{∞} the Galois \mathbf{P}^1 -bundle of) Theorem 1.2.

But if one traces through the construction of the Galois representation of Theorem 1.2 given in §1.2, one sees that this correspondence of deformations in (\ast_2) is precisely the Galois version of the *Kodaira-Spencer morphism* (reduced modulo p)

$$\mathbf{R}^1 f_{\mathrm{DR}, \ast}(\mathrm{Ad}(Q_{0,0} |_{\Delta_0})) \rightarrow \Omega_S = N_{\Delta_0/D_0^{\log}}$$

(where, in the notation of §1.1, we denote the structure morphism $X^{\log} \rightarrow S^{\log}$ by “ f ,” and the normal bundle of the diagonal Δ_0 in the mantle D_0^{\log} by “ $N_{\Delta_0/D_0^{\log}}$ ”) of the \mathbf{P}^1 -bundle $Q_{0,0}$ of the “formal \mathcal{MF}^{∇} -object” of Theorem 1.1. Note that this Kodaira-Spencer morphism is *equivariant* with respect to the natural action of Frobenius on both sides; thus, it is natural that *it* should be the “ \mathcal{MF}^{∇} -object version” of some Galois-theoretic object. For a more detailed discussion of this Kodaira-Spencer morphism in the “classical ordinary case,” we refer to [Mzk1], Chapter V, §1 (see, especially, the discussion of the classes “ η_{ε} ” and “ η^{Gal} ” in the subsections entitled “The Crystalline Induced \mathcal{MF}^{∇} -Object” and “Relation to the Canonical Affine Coordinates”).

At any rate, this Kodaira-Spencer morphism is clearly (cf. the theory of §1.1) *surjective modulo p* . This implies the *injectivity* of the correspondence of infinitesimal deformations of (\ast_2), as desired.

Thus, in order to complete the proof of Theorem 1.4, it remains to show that $\widehat{\kappa}_{S_{\mathbf{Q}_p}}$ is *injective on the fibers of*

$$\mathcal{N}[\infty]_{S_{\mathbf{Q}_p}} \rightarrow \tilde{Z}_{S_{\mathbf{Q}_p}}^{\log}$$

But this portion of the covering $\mathcal{N}[\infty] \rightarrow \mathcal{N}$ is precisely that portion which is the topic of the theory of *crystalline induction*. Thus, for this portion of the covering $\mathcal{N}[\infty] \rightarrow \mathcal{N}$, we may apply the theory of Theorem 1.2 (i.e., the theory of crystalline induction) *directly*. Note first that the profinite set corresponding to the covering $\mathcal{N}[\infty]_{S_{\mathbf{Q}_p}} \rightarrow \tilde{Z}_{S_{\mathbf{Q}_p}}^{\log}$ is precisely the set of \mathcal{O}_{ϖ} -valued points

$$\mathbf{A}_{\Phi}^{\text{Gal}}(\mathcal{O}_{\varpi})$$

of the object $\mathbf{A}_{\Phi}^{\text{Gal}}$ of Theorem 1.2. Moreover, the crystalline induction of Theorem 1.2 states precisely that the *morphism of profinite sets*

$$\mathbf{A}_{\Phi}^{\text{Gal}}(\mathcal{O}_{\varpi}) \rightarrow \mathbf{Rep}$$

(induced by the Galois \mathbf{P}^1 -bundle of Theorem 1.2) is *analytic* (cf. the Remark following Chapter IX, Theorem 2.11) relative to the given analytic structure of $\mathbf{A}_{\Phi}^{\text{Gal}}$ and the analytic structure of \mathbf{Rep} (over \mathcal{O}_{ϖ}) induced by the algebraic structure of $PGL_2(\mathcal{O}_{\varpi})$ by choosing topological generators of $\pi_1(\mathcal{C}_{\bar{\eta}})$ (cf. [Mzk1], Chapter V, §1, “Relation to the Canonical Affine Coordinates”). Moreover, the *derivative of this analytic morphism* is precisely the Galois version (this time over \mathbf{Z}_p , as opposed to modulo p) of the Kodaira-Spencer morphism

$$\mathbf{R}^1 f_{\text{DR},*}(\text{Ad}(Q_{0,0}|\Delta_0)) \rightarrow \Omega_S = N_{\Delta_0/D_0^{\log}}$$

considered above (cf. [Mzk1], Chapter V, §1, “Relation to the Canonical Affine Coordinates”). Thus, the derivative of the analytic morphism

$$\mathbf{A}_{\Phi}^{\text{Gal}}(\mathcal{O}_{\varpi}) \rightarrow \mathbf{Rep}$$

is *injective*. Put another way, this means that (noncanonically!) this morphism locally looks like the morphism

$$\begin{aligned} \text{Spf}(\mathcal{O}_{\varpi}[t_1, \dots, t_{3g-3+r}]^{\text{PD}, \wedge}) &\hookrightarrow \text{Spf}(\mathcal{O}_{\varpi}[t_1, \dots, \\ &t_{3g-3+r}, t_{(3g-3+r)+1}, \dots, t_{2(3g-3+r)}]^{\text{PD}, \wedge}) \end{aligned}$$

(where the t_i 's are indeterminates, and the superscripted “PD, \wedge ” stands for “the p -adic completion of the PD-envelope of the polynomial at the origin”). Thus, it is clear that this morphism is *injective* on the \mathcal{O}_{ϖ} -valued points of

$$\mathcal{O}_{\varpi}[t_1, \dots, t_{3g-3+r}]^{\text{PD}, \wedge}$$

i.e., the points obtained by sending t_1, \dots, t_{3g-3+r} to various elements of $p \cdot \mathcal{O}_\varpi$. This completes the proof of Theorem 1.4. \circ

Remark. The importance of Theorem 1.4 lies in the fact that it shows that the various p -adic analytic canonical indigenous bundles and Frobenius liftings constructed in the present work (cf., e.g., Chapter VII, Theorem 1.8) are *not isolated* analytic objects that have nothing to do with the “algebraic/arithmetic” outer Galois representation of $\pi_1((\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{Q}})$ on $\pi_1(\mathcal{C}_{\overline{\eta}})$, but rather may be regarded as a certain “portion” (cf. the term “immersion” in Theorem 1.4; Chapter VIII, Definition 1.4, Theorem 2.12) of the p -adic completion of the (algebraic!) coverings of $(\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{Q}_p}$ arising from this outer Galois representation – note that this is consistent with the “ \mathcal{MF}^∇ -object point of view” discussed in §1.3 of the Introduction to the present work.

Another way to regard Theorem 1.4 is as a partial analogue of the analysis in [KM] of *moduli of p -power torsion points on elliptic curves over \mathbf{Z}_p* . For instance, it is shown in [KM] that these moduli, which may be defined (just as \mathcal{R} above) as normalizations of natural log étale coverings over \mathbf{Q}_p , may also be constructed explicitly p -adically over the ordinary locus of the moduli stack of elliptic curves using Serre-Tate parameters. Thus, Theorem 1.4 may be interpreted as saying that the normalizations \mathcal{R} may be constructed p -adically over the Π -ordinary locus using the canonical Frobenius liftings of Chapter VII, Theorem 1.8. Note that (since the stacks discussed in [KM] may be regarded as moduli stacks of *isogenies* of elliptic curves) this point of view is consistent with the interpretation of “ \mathcal{Q}^Π ” as the stack of *quasi-analytic self-isogenies* discussed at the end of §1.5 of the Introduction to the present work.

Before proceeding, we make one more observation concerning Theorem 1.4 (in the present *binary ordinary* case). Namely,

Not only does the Π -ordinary theory allow one (i.e., by Theorem 1.4) to give a p -adic construction of a certain portion of $\widehat{\mathcal{R}}$ using indigenous bundles, it is, in fact, in some sense, the maximal such portion of $\widehat{\mathcal{R}}$.

The sense in which the Π -ordinary locus is the maximal locus on which one can p -adically construct $\widehat{\mathcal{R}}$ using indigenous bundles is the following:

The pair consisting of the tautological curve and its canonical Π -indigenous bundle over $B^+(\mathcal{N}) = B^+(\mathcal{N}[\infty])$ does not extend to any $B^+(U)$ for a larger p -adic open sub-object $U \supseteq \mathcal{N}[\infty]$ of $\widehat{\mathcal{R}}$.

(Here, we write $B^+(\mathcal{N}) = B^+(\mathcal{N}[\infty])$ in the sense that $B^+(-)$ of any of the finite subcoverings of $\mathcal{N}[\infty] \rightarrow \mathcal{N}$ whose limit is $\mathcal{N}[\infty] \rightarrow \mathcal{N}$ itself may be identified with $B^+(\mathcal{N})$.)

Indeed, if we base change via the natural morphisms (cf. the definition of $B^+(-)$, reviewed, for instance, at the beginning of Chapter IX, §2.3)

$$B^+(\mathcal{N}[\infty]) \rightarrow W((\mathcal{O}_{\tilde{\mathcal{N}}[\infty]} \otimes \mathbf{F}_p)_{\text{red}}); \quad B^+(U) \rightarrow W((\mathcal{O}_{\tilde{U}} \otimes \mathbf{F}_p)_{\text{red}})$$

(where $\tilde{\mathcal{N}}[\infty]$, \tilde{U} are the normalizations of $\mathcal{N}[\infty]$ and U , respectively, in their maximal log étale coverings in characteristic zero.), then one sees immediately that the existence of the resulting Π -indigenous bundle over $W((\mathcal{O}_{\tilde{U}} \otimes \mathbf{F}_p)_{\text{red}})$ contradicts the ω -closedness result of Chapter VII, Theorem 2.11. (Note that here, “ ω -closedness” implies closedness in the usual sense since $U_{\mathbf{F}_p}$ is open in $\widehat{\mathcal{R}}_{\mathbf{F}_p}$, which is an inverse limit of objects which are *finite* over $\mathcal{M}_{\mathbf{F}_p}$. Thus, the technical problems (cf., e.g., Chapter VII, Example 2.9) arising from infinite limits of blow-ups do not occur.)

§2. The Very Ordinary Spiked Case

Let k be a *finite field of odd characteristic p that contains \mathbf{F}_{p^2}* . Let $A = W(k)$; let K be the quotient field of A . Let S be a formally smooth, geometrically connected p -adic formal scheme over A of (constant) relative dimension d . Let us assume, moreover, that S is equipped with an étale morphism

$$S \rightarrow (\overline{\mathcal{N}}_{g,r}^{\Pi,s})_{\mathbf{Z}_p}^{\text{very ord}}$$

(for some *spiked VF-pattern Π of period 2*) into the *very ordinary locus* (cf. Chapter VII, Definition 3.6) of the associated shifted VF-stack. Let S^{\log} be the result of equipping S with the log structure pulled back from $\overline{\mathcal{M}}_{g,r}^{\log}$. We denote the divisor at infinity by $D \subseteq S$. Let

$$X^{\log} \rightarrow S^{\log}$$

be the pull-back of the *tautological log-curve*.

§2.1. The Formal \mathcal{MF}^∇ -Object

Let us denote by

$$\Phi^{\log} : S^{\log} \rightarrow S^{\log}$$

the *canonical Frobenius lifting* of Chapter VII, Theorem 1.8. Let $D \stackrel{\text{def}}{=} \text{Spf}(\mathcal{A})$. Thus, we have a *relative (p -adic formal) scheme with connection*

$$(D, \nabla_D)$$

over S^{\log} . Let

$$(F', \nabla_{F'})$$

be the *renormalized Frobenius pull-back* of (D, ∇_D) (cf. Chapter IX, Definition 1.10). As in Chapter IX, Theorem 3.1, we have a *Hodge subspace*

$$H \subseteq F'$$

On the other hand, in D , we have the *diagonal* $\Delta \subseteq D$ which serves as a “*Hodge subspace for D .*” As usual, we have *projection morphisms*

$$\beta, \gamma : D \rightarrow S$$

to the right and left factors, respectively.

Let \mathcal{P} be the *canonical indigenous bundle* on X^{\log} (cf. Chapter VII, Theorem 1.8). Now we would like to apply the *four step procedure* discussed in §1.1 to \mathcal{P} , i.e., we: (i.) *pull back* (X^{\log}, \mathcal{P}) via $\gamma : D \rightarrow S$; (ii.) *evaluate* the resulting crystal to obtain a \mathbf{P}^1 -bundle over $D \times_S X^{\log}$; and finally (iii) *push forward* by the projection $D \times_S X^{\log} \rightarrow X^{\log}$, so as to obtain an object

$$Q \rightarrow X$$

that is equipped with a natural *integrable connection* ∇_Q relative to $X^{\log} \rightarrow \text{Spec}(A)$. Moreover, (Q, ∇_Q) has the structure of an object – more precisely, a \mathbf{P}^1 -bundle – over

$$(D, \nabla_D) \times_S X$$

Finally, (Q, ∇_Q) is equipped with a *Hodge subspace* given by taking the image of the Hodge section $\Delta \rightarrow Q|_{\Delta}$ defined on the diagonal $\Delta \subseteq D$

(since over the diagonal, the \mathbf{P}^1 -bundle $(Q, \nabla_Q) \rightarrow (D, \nabla_D) \times_S X$ is indigenous).

By forming the *renormalized Frobenius pull-back* of (Q, ∇_Q) (cf. the discussion preceding Theorem 1.1), we obtain an object

$$(R, \nabla_R) \stackrel{\text{def}}{=} \mathbf{F}^*(\Phi_{S_{\mathbf{F}_p}}^*(Q, \nabla_Q) \times_{(\Phi_{S_{\mathbf{F}_p}}^* D)} F')$$

on X^{\log} which is equipped with an auxiliary structure of \mathbf{P}^1 -bundle over $(F', \nabla_{F'}) \times_S X$. Moreover, (R, ∇_R) is equipped with a *Hodge subspace* given by taking the image of the Hodge section

$$H \rightarrow R|_H$$

defined on $H \subseteq F'$. Note here that this Hodge section *exists* precisely because of the way that we *defined* H in the discussion preceding Chapter IX, Theorem 3.1 (cf. especially the definition of the subspace $H' \subseteq F'$ of *loc. cit.*).

Finally, by the “compatibility condition” of Chapter VII, Theorem 1.8, we have an isomorphism

$$(Q, \nabla_Q) \cong \mathbf{F}^*\{\Phi_{S_{\mathbf{F}_p}}^*(R, \nabla_R) \times_{(\Phi_{S_{\mathbf{F}_p}}^* F')} D\}$$

where the morphism

$$D \rightarrow \Phi_{S_{\mathbf{F}_p}}^* F'$$

implicit on the far right is that given by composing the *localization isomorphism* $D \rightarrow \mathbf{F}^*(F')$ (cf. Chapter IX, Lemma 1.15) with the *inclusion* $\mathbf{F}^*(F') \subseteq \Phi_{S_{\mathbf{F}_p}}^* F'$ (cf. the inclusion “ $F'' \subseteq \Phi_{S_{\mathbf{F}_p}}^* F'''$ ” of the discussion preceding Chapter IX, Lemma 1.14).

Note that this isomorphism defines a *Frobenius action* on the product of (Q, ∇_Q) and (R, ∇_R) (cf. Chapter IX, Theorem 1.16). In summary, we have the following result:

Theorem 2.1. *The object*

$$Q \times_X R$$

is equipped with all the structures necessary for a “formal” \mathcal{MF}^∇ -object on X^{\log} : full logarithmic connections (relative to $X^{\log} \rightarrow \text{Spec}(A)$), Hodge subspaces, and a

Frobenius action. Moreover, it is also equipped with a structure of \mathbf{P}^1 -bundle over the formal \mathcal{MF}^∇ -object

$$(D \times_S F') \times_S X$$

(where the implicit morphisms to S are the structure morphisms) on X^{\log} (cf. Chapter IX, Theorem 1.16). Finally, this structure of \mathbf{P}^1 -bundle is compatible with the various full logarithmic connections, Hodge subspaces, and Frobenius actions involved.

§2.2. The Crystalline Induced Representation

In this subsection, we carry out the *spiked analogue* of the construction discussed in §1.2.

Let

$$Z_K^{\log} \rightarrow S_K^{\log}$$

be as in the discussion concerning the *spiked Galois mantle* preceding Chapter IX, Definition 2.12. As in §1.2, we have *morphisms* (cf. the discussion following Chapter IX, Definition 2.12)

$$\mathcal{A}(B^+(S^{\log})) \rightarrow (\mathcal{A}_\Phi^{\text{Gal}})_{B^+}$$

and

$$\mathcal{F}'(B^+(S^{\log})) \rightarrow (\mathcal{A}_\Phi^{\text{Gal}})_{B^+}^F$$

(where $(\mathcal{A}_\Phi^{\text{Gal}})_{B^+} \stackrel{\text{def}}{=} B^+(S^{\log}) \otimes_{W(\mathbf{F}_{p^2})} \mathcal{A}_\Phi^{\text{Gal}}$) that *respect Hodge filtrations, Frobenius, and the action of $\Pi_{Z^{\log}}$* . By *base-changing* (as in §1.2) the objects obtained by evaluating the \mathbf{P}^1 -bundles (Q, ∇_Q) and (R, ∇_R) on $B^+(U^{\log})$ (for a “small” p -adic open $U^{\log} \subseteq X^{\log}$), by means of these morphisms, and taking Frobenius invariants as in §1.2, we thus obtain the following result (cf. [Mzk1], Chapter V, Theorem 1.4):

Theorem 2.2. *Let*

$$\mathcal{A}_\Phi^{\text{Gal}}$$

be the total Galois mantle (cf. Chapter IX, Definition 2.12) associated to the canonical very ordinary spiked Frobenius lifting Φ^{\log} and Hodge subspace H under consideration (cf. Chapter IX, Theorem 3.1). Let

$$\mathbf{A}_{\Phi}^{\text{Gal}} \stackrel{\text{def}}{=} \text{Spf}(\mathcal{A}_{\Phi}^{\text{Gal}})$$

(where “ Spf ” is taken with respect to the p -adic topology). Then associated to the formal \mathcal{MF}^{∇} -object of Theorem 2.1, we have a \mathbf{P}^1 -bundle

$$\mathbf{P}_{\Phi}^{\text{Gal}} \rightarrow \mathbf{A}_{\Phi}^{\text{Gal}}$$

equipped with a continuous action by $\pi_1((X_Z^{\log})_K)$ (where $X_Z^{\log} \stackrel{\text{def}}{=} X^{\log} \times_{S^{\log}} Z^{\log}$) which is compatible with the action of $\Pi_{Z^{\log}} \stackrel{\text{def}}{=} \pi_1(Z_K^{\log})$ on $\mathbf{A}_{\Phi}^{\text{Gal}}$.

Remark. Just as was the case for Theorem 1.2, once one has this Galois \mathbf{P}^1 -bundle $\mathbf{P}_{\Phi}^{\text{Gal}}$, one can carry out a discussion of the relationship between it and the various types of (strong and weak) canonical parameters of Chapter IX in a fashion entirely analogous to that of the discussion preceding [Mzk1], Chapter V, Theorem 1.7. (In particular, one obtains a spiked analogue of the observation (*) immediately following [Mzk1], Chapter V, Theorem 1.7.) Since the proofs and results are entirely similar to those of [Mzk1], Chapter V, §1, we leave it to the reader to make them explicit.

§2.3. Relation to the Profinite Teichmüller Group

In this subsection, we use the notation of §1.4. Our goal in this subsection is to discuss the *spiked analogue* of §1.4, i.e., to relate the various p -adic analytic objects constructed in the spiked case to the algebraic/analytic outer Galois action of $\pi_1((\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{Q}})$ on $\pi_1(\mathcal{C}_{\overline{\eta}})$.

Throughout this subsection, the period ϖ will be 2, i.e.:

$$\varpi = 2$$

In particular, we have $\mathcal{O}_{\varpi} = W(\mathbf{F}_{p^2})$. Then, just as in §1.4, we have a continuous action of $\pi_1(\mathcal{M}_{\mathbf{Q}_p})$ on

$$\mathbf{Rep} \stackrel{\text{def}}{=} \text{Rep}(\pi_1(\mathcal{C}_{\overline{\eta}}), \text{PGL}_2(\mathcal{O}_{\varpi}))$$

This gives rise to an infinite covering

$$\mathcal{R}_{\mathbf{Q}_p} \rightarrow \mathcal{M}_{\mathbf{Q}_p}$$

whose normalization we denote by \mathcal{R} .

Next, let us write

$$\mathcal{N} \stackrel{\text{def}}{=} (\overline{\mathcal{N}}_{g,r}^{\Pi,s})_{\mathbf{Z}_p}^{\text{very ord, log}}$$

Recall that we have a natural morphism $\mathcal{N} \rightarrow \mathcal{M}$. Next, recall the “set-theoretic Galois representation” \mathcal{V} of Chapter IX, Definition 2.5. This set \mathcal{V} is a profinite set with a continuous $\pi_1(\mathcal{N}_{\mathbf{Q}_p})$ -action. Thus, it corresponds to some *infinite covering* (i.e., inverse limit of finite log étale coverings)

$$\mathcal{N}[\infty]_{\mathbf{Q}_p} \rightarrow \mathcal{N}_{\mathbf{Q}_p}$$

We will denote the normalization of \mathcal{N} in $\mathcal{N}[\infty]_{\mathbf{Q}_p}$ by

$$\mathcal{N}[\infty] \rightarrow \mathcal{N}$$

It follows immediately from the definitions that $Z_{\mathbf{Q}_p}^{\log}$ (where Z^{\log} is as in Theorem 2.2 in the case where we take “ S^{\log} ” to be \mathcal{N} itself) is an *intermediate covering* between $\mathcal{N}_{\mathbf{Q}_p}$ and $\mathcal{N}[\infty]_{\mathbf{Q}_p}$. Let us write \tilde{Z}^{\log} for the *normalization* of \mathcal{N} in $Z_{\mathbf{Q}_p}^{\log}$. Thus, we have morphisms

$$\pi_1(\mathcal{N}[\infty]_{\mathbf{Q}_p}) \rightarrow \pi_1(Z_{\mathbf{Q}_p}^{\log}) \rightarrow \pi_1(\mathcal{N}_{\mathbf{Q}_p})$$

In particular, we have a natural action of $\pi_1(\mathcal{N}[\infty]_{\mathbf{Q}_p})$ on $\mathbf{A}_{\Phi}^{\text{Gal}}$ (cf. Theorem 2.2). Moreover, $\mathbf{A}_{\Phi}^{\text{Gal}}$ admits a *tautological* \mathcal{O}_{ϖ} -valued point

$$\alpha_{\infty} \in \mathbf{A}_{\Phi}^{\text{Gal}}(\mathcal{O}_{\varpi})$$

which is *fixed* by $\pi_1(\mathcal{N}[\infty]_{\mathbf{Q}_p})$ (cf. Chapter IX, Corollary 2.6). Thus, if one restricts the “Galois \mathbf{P}^1 -bundle” of Theorem 2.2 to the point α_{∞} (and uses the fact that the automorphisms of \mathbf{P}^1 are given schematically by “ PGL_2 ”), we obtain from the definition of \mathbf{Rep} that we have a natural *classifying morphism*

$$\kappa_{\mathbf{Q}_p} : \mathcal{N}[\infty]_{\mathbf{Q}_p} \rightarrow \mathcal{R}_{\mathbf{Q}_p}$$

for the restriction to α_{∞} of the “Galois \mathbf{P}^1 -bundle” of Theorem 2.2. Moreover, it follows from the definition (as normalizations) of $\mathcal{N}[\infty]$ and \mathcal{R} that $\kappa_{\mathbf{Q}_p}$ extends to a morphism

$$\kappa : \mathcal{N}[\infty] \rightarrow \mathcal{R}$$

Write $\widehat{\mathcal{M}}$ for the *p-adic completion* $\widehat{\mathcal{M}}$ of \mathcal{M} . Since \mathcal{M} is *proper* over \mathbf{Z}_p , it follows that finite log étale coverings of $\widehat{\mathcal{M}}_{\mathbf{Q}_p}$ are *equivalent* to

finite log étale coverings of $\mathcal{M}_{\mathbf{Q}_p}$. Let us write $\widehat{\mathcal{R}} \stackrel{\text{def}}{=} \mathcal{R} \times_{\mathcal{M}} \widehat{\mathcal{M}}$. (Note: $\widehat{\mathcal{R}}$ is a *limit* of p -adically complete finite coverings of $\widehat{\mathcal{M}}$, not a p -adic completion of a *limit of finite coverings*. In particular, (in general) $\widehat{\mathcal{R}}$ is not p -adically complete.) Thus, it follows from the definitions that κ defines morphisms

$$\widehat{\kappa} : \mathcal{N}[\infty] \rightarrow \widehat{\mathcal{R}}; \quad \widehat{\kappa}_{\mathbf{Q}_p} : \mathcal{N}[\infty]_{\mathbf{Q}_p} \rightarrow \widehat{\mathcal{R}}_{\mathbf{Q}_p}$$

The purpose of the present subsection is to prove the following result:

Theorem 2.3. *We have a natural commutative diagram:*

$$\begin{array}{ccc} \mathcal{N}[\infty] & \xrightarrow{\widehat{\kappa}} & \widehat{\mathcal{R}} \\ \downarrow & & \downarrow \\ \mathcal{N} & \longrightarrow & \widehat{\mathcal{M}} \end{array}$$

Moreover, the morphism $\widehat{\kappa}$ – which is induced by the classifying morphism for the restriction to the tautological point $\alpha_{\infty} \in \mathbf{A}_{\Phi}^{\text{Gal}}(\mathcal{O}_{\varpi})$ of the “Galois \mathbf{P}^1 -bundle” of Theorem 2.2 – is an open immersion.

Proof. The proof follows the precisely the same pattern – i.e., involving *three levels of distinction* – as the proof of Theorem 1.4. In fact, we presented the proof of Theorem 1.4 in such a way that this would be the case. The careful reader will observe that, in the binary case of Theorem 1.4, because of the affine linear structure on the total Galois mantle, it really is not so necessary to distinguish between the second and third levels of distinction. In the present spiked case, by contrast, because the total Galois mantle *does not carry an affine linear structure*, the argument for the “second level of distinction” must be distinguished from the argument for the “third level of distinction” (since the *linear* structure of deformations used in the argument for the “second level of distinction” only exists modulo p). In the argument for the “second level of distinction” in the present spiked case, we use the interpretation of $Z_{\mathbf{Q}_p}^{\log}$ (respectively, $T_{\mathbf{Q}_p}^{\log}$) – as parametrizing splittings of a certain Galois torsor – given in the discussion preceding Chapter IX, Definition 2.12 (respectively, following Chapter IX, Theorem 1.5). In the argument for the “third level of distinction” in the present spiked case, we use the theory of crystalline induction (in the spiked case) summarized in Theorem 2.2. \circ

Remark. Just as was the case for Theorem 1.4, the importance of Theorem 2.3 lies in the fact that it (together with Chapter IX, Corollary 2.6) shows that the various p -adic analytic canonical indigenous bundles and

Frobenius liftings constructed in the spiked case (cf., e.g., Chapter VII, Theorem 1.8) are *not isolated* analytic objects that have nothing to do with the “algebraic/arithmetic” *outer Galois representation* of $\pi_1((\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{Q}})$ on $\pi_1(\mathcal{C}_{\overline{\eta}})$, but rather may be regarded as a certain “portion” (cf. the term “immersion” in Theorem 2.3) of the *p-adic completion* of the (algebraic!) coverings of $(\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{Q}_p}$ arising from this outer Galois representation – cf. also the discussions of §1.3, 1.5, of the Introduction to the present work, as well as the Remark following Theorem 1.4.

§3. Conclusion

Finally, we pause to take a look at what we have achieved. Just as in the Introduction to the present work, we would like to *describe the p-adic theory by comparing it to the classical theory at the infinite prime*. Thus, let us write

$$\mathcal{C}_{\mathbf{C}} \rightarrow \mathcal{M}_{\mathbf{C}} \stackrel{\text{def}}{=} (\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{C}}$$

for the *universal log-curve over the moduli stack* $(\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{C}}$ of *r*-pointed stable log-curves of genus *g* over the complex numbers. Let us fix a “base-point” (say, in the interior – i.e., the open sub-object where the log structure is trivial – of $\mathcal{M}_{\mathbf{C}}$)

$$[X] \in \mathcal{M}_{\mathbf{C}}(\mathbf{C})$$

corresponding to some hyperbolic algebraic curve *X* over \mathbf{C} . Let us write

$$\mathcal{X} \stackrel{\text{def}}{=} X(\mathbf{C})$$

for the corresponding *hyperbolic Riemann surface*. Next, let us consider the space

$$\mathbf{Rep}_{\mathbf{C}} \stackrel{\text{def}}{=} \text{Rep}(\pi_1^{\text{top}}(\mathcal{X}), PGL_2(\mathbf{C}))$$

of isomorphism classes of representations of the *topological fundamental group* $\pi_1^{\text{top}}(\mathcal{X})$ into $PGL_2(\mathbf{C})$. This space has the structure of an *algebraic variety* over \mathbf{C} , induced by the algebraic structure of $PGL_2(\mathbf{C})$ by choosing generators of $\pi_1^{\text{top}}(\mathcal{X})$. Note, moreover, that as $[X]$ varies, the resulting spaces $\text{Rep}(\pi_1(\mathcal{X}), PGL_2(\mathbf{C}))$ form a *local system* on $\mathcal{M}_{\mathbf{C}}$ (valued in the category of algebraic varieties over \mathbf{C}) which we denote by

$$\mathcal{R}_{\mathbf{C}} \rightarrow \mathcal{M}_{\mathbf{C}}$$

One can also think of $\mathcal{R}_{\mathbf{C}}$ as the local system defined by the natural action of $\pi_1^{\text{top}}(\mathcal{M}_{\mathbf{C}}(\mathbf{C}))$ on


$$\text{Rep}_{\mathbf{C}} \stackrel{\text{def}}{=} \text{Rep}(\pi_1^{\text{top}}(\mathcal{X}), \text{PGL}_2(\mathbf{C}))$$

(induced by the *natural outer action* of $\pi_1^{\text{top}}(\mathcal{M}_{\mathbf{C}}(\mathbf{C}))$ on $\pi_1^{\text{top}}(\mathcal{X})$) – cf. the discussion of the p -adic case in §1.4, 2.3.

Next, let us denote by

$$\mathcal{QF} \subseteq \mathcal{R}_{\mathbf{C}}$$

the subspace whose fiber over a point $[Y] \in \mathcal{M}_{\mathbf{C}}(\mathbf{C})$ is given by the representations $\pi_1^{\text{top}}(\mathcal{Y}) \rightarrow \text{PGL}_2(\mathbf{C})$ that define *quasi-Fuchsian groups* (cf. Introduction, §0.4), i.e., simultaneous uniformizations of pairs of Riemann surfaces (of the same type (g, r)), for which one (say, the “first” one) of the pair of Riemann surfaces uniformized is the Riemann surface \mathcal{Y} corresponding to $[Y]$. Thus, whereas the fibers of $\mathcal{R}_{\mathbf{C}} \rightarrow \mathcal{M}_{\mathbf{C}}$ are of *dimension* $2(3g - 3 + r)$ over \mathbf{C} , the fibers of $\mathcal{QF} \rightarrow \mathcal{M}_{\mathbf{C}}$ are of *dimension* $3g - 3 + r$ over \mathbf{C} .



Then, relative to the notation of §1.4, 2.3, the *analogy between the complex and p -adic cases* may be summarized by the following diagram:

$$\begin{array}{ccc} \mathcal{QF} & \Longleftrightarrow & \mathcal{N}[\infty] \\ \cap & & \cap \\ (\mathcal{R}_C \rightarrow \mathcal{M}_C) & \Longleftrightarrow & (\widehat{\mathcal{R}} \rightarrow \widehat{\mathcal{M}}) \end{array}$$

(where the vertical inclusion on the left is the natural one; and the vertical inclusion on the right is the morphism $\widehat{\kappa}$ of Theorems 1.4, 2.3). We also give an *illustration* above of this sort of situation. Relative to this illustration, the “integral subspaces” of the local system are \mathcal{QF} and $\mathcal{N}[\infty]$ (cf. §0.4 of the Introduction of the present work for an explanation of the term “integral”). Note that just as in the complex case, the fibers of the covering $\mathcal{N}[\infty] \rightarrow \widehat{\mathcal{M}}$ have, so to speak, “Galois dimension” $3g - 3 + r$ over \mathcal{O}_{ϖ} (cf. the crystalline induction portion of the proof of Theorem 1.4), whereas the fibers of the covering $\widehat{\mathcal{R}} \rightarrow \widehat{\mathcal{M}}$ are of “Galois dimension” $2(3g - 3 + r)$ over \mathcal{O}_{ϖ} . In the p -adic case, $\mathcal{N}[\infty]$ denotes the “*crystalline*” or “*Frobenius invariant indigenous bundle*” locus of $\widehat{\mathcal{R}}$ – cf. the discussion of §0.4 of the Introduction.

In the complex case, the “Frobenius” (i.e., complex conjugation) invariant portion of \mathcal{QF} is the space of *Fuchsian groups*, hence defines the *Bers uniformization* of Teichmüller space (cf. §0.5 of the Introduction). On the other hand, in the p -adic case, the covering $\mathcal{N}[\infty] \rightarrow \widehat{\mathcal{M}}$ is “made up of” composites of Frobenius liftings, by *forgetting* that these Frobenius liftings are morphisms from a single space to itself, and just thinking of them as *coverings*. If one then invokes the structure of Frobenius liftings as *morphisms from a single space to itself*, one so-to-speak recovers the original Frobenius liftings, which (by the theory of Chapters VIII and IX) define *p -adic uniformizations* of $(\overline{\mathcal{M}}_{g,r}^{\log})_{\mathbf{Z}_p}$.

In the complex case, the space of quasi-Fuchsian groups \mathcal{QF} may also be interpreted in terms of *quasi-conformal maps*. Similarly, in the p -adic case, one may interpret integral Frobenius invariant indigenous bundles as *quasi-analytic self-isogenies* of hyperbolic curves (cf. the end of §1.5 of the Introduction).

Finally, in the complex case, although \mathcal{QF} is *not closed* in \mathcal{R}_C , the space \mathcal{QF} (when regarded as a space of representations) is *complete* relative to the condition that the representations always define *indigenous bundles* for some conformal structures on the *two* surfaces being uniformized. Note that one may think of these two surfaces as *reflections* of another, i.e., translates of one another by some action of *Frobenius* at the infinite prime (i.e., complex conjugation). Similarly, although $\mathcal{N}[\infty]$ is *not closed* in $\widehat{\mathcal{R}}$, it is *complete* (at least for binary VF-patterns Π) in the sense discussed at the end of §1.4, i.e., relative to the condition that the representation always defines an indigenous bundle on

the universal thickening $B^+(-)$ of the base. Note that this thickening $B^+(-)$ is in some sense the *minimal thickening* of (the normalization of the maximal log étale in characteristic zero extension of) “ $(-)$ ” that admits an action of Frobenius (cf. the theory of Chapter VI, §1; $B^+(-)$ is the PD-completion of the rings $B(-)$ discussed in Chapter VI, §1; in fact, instead of using $B^+(-)$ here, it would also be quite sufficient to use the rings $B(-)$ of Chapter VI, §1). In other words, just as in the complex case,

$\mathcal{N}[\infty]$ is already complete relative to the condition that the representations it parametrizes always define indigenous bundles on the given curve and all of its Frobenius conjugates.

Appendix: Ordinary Stable Bundles on a Curve

§0. Introduction

The purpose of this Appendix is to discuss another example of *canonical Frobenius actions* and the analogy between such Frobenius actions (at finite primes) and Kähler metrics at the infinite prime (cf. §0.8 of the Introduction to the present work). First, let us recall that:

Given a *curve in positive characteristic*, it is natural to study the curve by considering the natural *Frobenius action* induced on various objects canonically associated to the curve.

Well-known examples of this include the Frobenius action on *rational points* of the curve, the Frobenius action on its *Jacobian*, and the Frobenius action on its *étale or crystalline cohomology*. The main body of the present work is concerned, by contrast, with the natural Frobenius action on the *moduli space of hyperbolic curves equipped with an indigenous bundle*. In this Appendix, we would like to discuss the Frobenius action on the *moduli space of stable vector bundles* on a curve.

We begin by reviewing the algebraic and complex theories, and then proceed to discuss the (ordinary) p -adic theory. One aspect of the case of moduli of stable bundles (as opposed to moduli of curves equipped with an indigenous bundle, as treated in [Mzk1] and the present work) is that it is technically *much simpler*, and yet serves to illustrate what the author believes to be essentially the same phenomenon, namely the analogy between *Kähler metrics* (in the complex case) and *Frobenius actions* (in the p -adic case). On the other hand, one important difference between the p -adic Frobenius actions in the stable bundle case, on the one hand, and the cases of moduli of abelian varieties (Serre-Tate theory) and moduli of curves (the theory of [Mzk1]

and the present work) is that in the former case, the action of Frobenius on the cotangent space of the space in question has *slope zero*, whereas in the latter case, the action has *positive slope*. This difference makes a big difference in the way one constructs canonical coordinates, and, in particular, in the sort of base over which the canonical coordinates are defined. For instance, in the stable bundle case, one has to limit oneself to the Witt vectors of a *finite field* and use relatively strong results on the eigenvalues of Frobenius – namely, the *Weil Conjectures* – in order to construct the canonical coordinates. By contrast, for moduli of curves and abelian varieties, the base can be the Witt vectors of any perfect scheme, and one does not need any deep results on the eigenvalues of Frobenius in order to construct the canonical coordinates – that is, the *positivity* of the slopes does all the work.

We remark that there is little that is original in this Appendix: §1, 2 are classical and well-known, and §3 is an immediate consequence of results which are classical and well-known. (Indeed, in some sense, one may regard the theory of §3 as a rather trivial special case of the theory of the category \mathcal{MF}^∇ discussed in [Falt1], §2.) After reviewing basic algebraic definitions and results in §1, we go on in §2, 3 to discuss the following two types of uniformization result:

- (1) showing that “any” stable bundle can be constructed from a *representation of the fundamental group* of the curve into some compact subgroup of GL_r (Theorems 2.1 and 3.12);
- (2) showing the existence of *canonical local affine uniformizations* of the moduli space of stable bundles that generalize the exponential map in the abelian case (Theorems 2.5 and 3.16).

Relative to this point of view, §2 discusses these sorts of results at infinite primes, while §3 discusses these sorts of results at finite primes.

§1. The Algebraic Theory

We begin by reviewing the basic non-arithmetic theory.

§1.1. Basic Definitions

Let S be a scheme. Let

$$f : X \rightarrow S$$

be a *curve*, by which we shall mean (in this Appendix) that f is smooth, proper, geometrically connected, and of relative dimension one. Let \mathcal{E} be a *vector bundle* on X of rank r , where r is *invertible* on S . Let

$$\mathrm{Ad}(\mathcal{E}) \subseteq \mathrm{End}(\mathcal{E})$$

be the subbundle of trace zero endomorphisms. We shall say that \mathcal{E} is *special* if its determinant is trivial.

Definition 1.1. Suppose that \mathcal{E} is special. If $S = \mathrm{Spec}(k)$, where k is an algebraically closed field, then we shall say that \mathcal{E} is *stable* if every nonzero proper (i.e., $\neq \mathcal{E}$) subsheaf of \mathcal{E} has negative degree. For arbitrary S , we shall say that \mathcal{E} is *stable* if its restriction to every geometric fiber of f is stable.

Suppose that \mathcal{E} is special, and that $\nabla_{\mathcal{E}}$ is a connection on \mathcal{E} relative to $f: X \rightarrow S$. We shall say that $\nabla_{\mathcal{E}}$ is *special* if the connection it induces on $\det(\mathcal{E}) = \mathcal{O}_X$ is the trivial one.

Proposition 1.2. *Suppose that \mathcal{E} is special and stable. Then the (fine moduli) space of special connections on \mathcal{E} forms a (nonempty) torsor on S over the vector bundle*

$$f_*(\omega_{X/S} \otimes_{\mathcal{O}_X} \mathrm{Ad}(\mathcal{E}))$$

Moreover, if $X \rightarrow S$ is a curve of genus g , then this vector bundle is of rank $(r^2 - 1)(g - 1)$.

Proof. Locally on X , \mathcal{E} is trivial, hence admits a special connection. The difference between two special connections forms a section of $\omega_{X/S} \otimes \mathrm{Ad}(\mathcal{E})$. Moreover, $\mathbf{R}^1 f_*(\omega_{X/S} \otimes \mathrm{Ad}(\mathcal{E}))$ is dual to $f_* \mathrm{Ad}(\mathcal{E})$ (since $\mathrm{Ad}(\mathcal{E})$ is self-dual – here, we use that r is invertible on S). Thus, since the obstruction to putting a special connection on \mathcal{E} forms a section of $\mathbf{R}^1 f_*(\omega_{X/S} \otimes \mathrm{Ad}(\mathcal{E}))$, it suffices to show that $f_* \mathrm{Ad}(\mathcal{E}) = 0$. (Note that the rank of the vector bundle $f_*(\omega_{X/S} \otimes_{\mathcal{O}_X} \mathrm{Ad}(\mathcal{E}))$ may be computed by Riemann-Roch.)

Let us prove that $f_* \mathrm{Ad}(\mathcal{E}) = 0$. Note that it suffices to prove this when $S = \mathrm{Spec}(k)$, where k is an algebraically closed field k . Then if $\mathcal{E} \rightarrow \mathcal{E}$ is a morphism, write \mathcal{I} for its *image*. If $\mathcal{I} \neq 0$ or \mathcal{E} , then Definition 1.1 implies that its degree is both *positive* (since \mathcal{I} is a quotient of \mathcal{E}) and *negative* (since \mathcal{I} includes into \mathcal{E}). This is absurd. Thus, \mathcal{I} is zero or the whole of \mathcal{E} . This proves that $\mathrm{End}_{\mathcal{O}_X}(\mathcal{E})$ is a division algebra of finite dimension over k . Since k is algebraically closed, this implies that $\mathrm{End}_{\mathcal{O}_X}(\mathcal{E}) = k$, so $f_* \mathrm{Ad}(\mathcal{E}) = 0$, as desired. \circ

§1.2. Moduli

We have the following result (cf., for instance, Théorème 17, p. 22, of [Sesh]):

Theorem 1.3. *The coarse moduli space of rank r special stable bundles on X is represented by a smooth, quasi-compact, (relative) algebraic space*

$$\mathcal{M}_r(X/S) \rightarrow S$$

The coarse moduli space of rank r special stable bundles on X equipped with a special connection is represented by a smooth, quasi-compact, (relative) algebraic space

$$\mathcal{A}_r(X/S) \rightarrow S$$

which has an additional structure of affine torsor over the sheaf of differentials $\Omega_{\mathcal{M}_r(X/S)/S}$ on $\mathcal{M}_r(X/S)$.

Proof. Théorème 17, p. 22, of [Sesh] deals with $\mathcal{M}_r(X/S)$ when S is the spectrum of an algebraically closed field. The general case is not difficult, since we only need here that $\mathcal{M}_r(X/S)$ and $\mathcal{A}_r(X/S)$ are algebraic spaces (not necessarily quasi-projective S -schemes), and has been considered by numerous authors. One approach is simply to mimic the proof of Theorem 2.7 of Chapter I. (Note that since we are working with vector bundles *whose determinant is trivial*, it follows that the automorphism group of \mathcal{E} is always *finite*.) The assertions on $\mathcal{A}_r(X/S) \rightarrow S$ follow immediately from Proposition 1.2 (and the fact that the vector bundle of Proposition 1.2 may be identified with the sheaf of differentials $\Omega_{\mathcal{M}_r(X/S)/S}$). \circ

Proposition 1.4. *Suppose that $(\mathcal{E}, \nabla_{\mathcal{E}})$ is a special, stable bundle with special connection on X . Then the tangent bundle to $\mathcal{M}_r(X/S)$ (respectively, $\mathcal{A}_r(X/S)$) at the point defined by \mathcal{E} (respectively, $(\mathcal{E}, \nabla_{\mathcal{E}})$) is given by $\mathbf{R}^1 f_* \mathrm{Ad}(\mathcal{E})$ (respectively, $\mathbf{R}^1 f_{\mathrm{DR},*} \mathrm{Ad}(\mathcal{E}, \nabla_{\mathcal{E}})$). Moreover, the de Rham cohomology fits into a self-dual exact sequence*

$$0 \rightarrow f_*(\omega_{X/S} \otimes \mathrm{Ad}(\mathcal{E})) \rightarrow \mathbf{R}^1 f_{\mathrm{DR},*} \mathrm{Ad}(\mathcal{E}, \nabla_{\mathcal{E}}) \rightarrow \mathbf{R}^1 f_* \mathrm{Ad}(\mathcal{E}) \rightarrow 0$$

given by subquotients of the Hodge filtration on the de Rham cohomology. Finally, the surjection on the right (of the above exact sequence) may be identified with the morphism on tangent bundles (at $(\mathcal{E}, \nabla_{\mathcal{E}})$) induced by the natural projection $\mathcal{A}_r(X/S) \rightarrow \mathcal{M}_r(X/S)$.

§2. The Complex Theory

The main results here concern two types of *uniformizations* of $\mathcal{M}_r(X/S)$: one given by regarding stable bundles as obtained from *unitary representations of the fundamental group* of the curve; the other obtained by considering the *canonical coordinates* associated to a certain natural real analytic Kähler metric on $\mathcal{M}_r(X/S)$. In this §, we suppose that $S = \text{Spec}(\mathbf{C})$, and we write

$$\Pi \stackrel{\text{def}}{=} \pi_1^{\text{an}}(X)$$

for the *analytic fundamental group* $\pi_1^{\text{an}}(X)$ of X (with respect to some base-point). Also, we shall write

$$\tilde{X} \rightarrow X$$

for the (analytic) universal covering space of X .

§2.1. Unitary Representations of the Fundamental Group

Let V be a finite dimensional complex vector space. Then given a representation

$$\rho : \Pi \rightarrow \text{GL}(V)$$

we can construct a *vector bundle*

$$\mathcal{E}_\rho$$

on X by taking the quotient of $\tilde{X} \times V$ by Π (which acts on both factors – on the first factor “by nature,” and on the second factor by means of ρ). Thus, we have defined a way of associating vector bundles to representations of Π . Relative to this correspondence, we have the following result (Théorème 39, p. 41, of [Sesh]):

Theorem 2.1. *Isomorphism classes of special, stable bundles on X of rank r are in one-to-one correspondence with isomorphism classes of irreducible unitary representations*

$$\rho : \Pi \rightarrow \text{SU}_r$$

(whose determinant is trivial).

Here,

$$\mathrm{SU}_r \subseteq \mathrm{SL}_r(\mathbf{C})$$

is the group of *special unitary matrices*. Let us write (for simplicity)

$$\mathcal{M}_r(X) \stackrel{\mathrm{def}}{=} \mathcal{M}_r(X/\mathbf{C})(\mathbf{C})$$

for the *complex manifold* given by the \mathbf{C} -valued points $\mathcal{M}_r(X/\mathbf{C})(\mathbf{C})$. Then, in terms of moduli, we can reinterpret Theorem 2.1 as saying that we have a sort of *real analytic uniformization* of $\mathcal{M}_r(X)$:

Corollary 2.2. *We have a real analytic isomorphism*

$$\mathcal{M}_r(X/\mathbf{C})(\mathbf{C}) \cong \mathrm{Hom}^{\mathrm{irred}}(\Pi, \mathrm{SU}_r)/\sim$$

where “ \sim ” stands for “up to isomorphism of representations,” and the superscripted “irred” stands for “irreducible representations.”

Note that if \mathcal{E}_ρ arises from a ρ that maps into SU_r , it follows that \mathcal{E}_ρ gets a *natural Hermitian metric* $\langle -, - \rangle_\rho$. (Indeed, the fact that ρ maps into SU_r implies that the vector space V admits a Hermitian metric stabilized by the action of Π ; thus, by descent (from \tilde{X} to X), we obtain a metric on \mathcal{E}_ρ .) By Theorem 2.1, every special, stable \mathcal{E} arises from a *unique* such ρ (up to isomorphism). Thus, we see in particular that:

Corollary 2.3. *Every special, stable \mathcal{E} on X admits a natural Hermitian metric $\langle -, - \rangle_\mathcal{E}$ whose dependence on \mathcal{E} is real analytic.*

Finally, note that whether ρ is unitary or not, every vector bundle \mathcal{E}_ρ gets a natural holomorphic connection ∇_ρ obtained by taking the sections of $\mathcal{E}|_{\tilde{X}} = \tilde{X} \times V$ arising from V to be *constant*. Thus, in particular, if \mathcal{E} is special and stable, Theorem 2.1 allows us to write \mathcal{E} as \mathcal{E}_ρ for a uniquely defined unitary representation ρ , hence gives us a *canonical choice of connection* $\nabla_\mathcal{E}$ on \mathcal{E} . That is to say, we have the following:

Corollary 2.4. *Every special, stable \mathcal{E} on X admits a natural $\nabla_\mathcal{E}$ whose dependence on \mathcal{E} is real analytic. In other words, the (nontrivial) holomorphic affine torsor $\mathcal{A}_r(X) \rightarrow \mathcal{M}_r(X)$ admits a natural real analytic section*

$$s : \mathcal{M}_r(X) \rightarrow \mathcal{A}_r(X)$$

(given by $\mathcal{E} \mapsto (\mathcal{E}, \nabla_\mathcal{E})$).

§2.2. The Kähler Approach

One aspect of the uniformization of Corollary 2.2 is that it still does not give us *canonical coordinates* on $\mathcal{M}_r(X)$, even locally. One way to view this problem is as follows: In the “abelian case,” i.e., when one works with line bundles rather than special vector bundles, one has the morphism

$$\dots \longrightarrow H^1(X, \mathcal{O}_X) \xrightarrow{\exp} H^1(X, \mathcal{O}_X^\times) = \text{Pic}(X) \longrightarrow \dots$$

induced by the exponential exact sequence on X . That is to say, one has a uniformization of $\text{Pic}(X)$ by the affine space $H^1(X, \mathcal{O}_X)$. It is thus natural to wonder if one can find an analogous morphism when one replaces “ \mathbf{G}_m ” by “ $\text{SL}_r(-)$.” Unfortunately, in the non-abelian case, things are not so simple, and in order to find such canonical affine coordinates (even locally), one must introduce some *geometry*, as follows.

Let \mathcal{E} be a *special stable vector bundle* on X . Let

$$V_{\mathcal{E}} = H^0(X, \omega_X \otimes_{\mathcal{O}_X} \text{Ad}(\mathcal{E}))$$

Then let us note that the canonical metric $\langle -, - \rangle_{\mathcal{E}}$ (cf. Corollary 2.3) on \mathcal{E} naturally defines a *canonical Hermitian metric* $\langle -, - \rangle_{V_{\mathcal{E}}}$ on the complex vector space $V_{\mathcal{E}}$, as follows: If $\theta, \psi \in V_{\mathcal{E}}$, then $\langle \theta, \psi \rangle_{\mathcal{E}}$ may be regarded as a $(1, 1)$ -form on X ; then we define

$$\langle \theta, \psi \rangle_{V_{\mathcal{E}}} \stackrel{\text{def}}{=} \int_X \langle \theta, \psi \rangle_{\mathcal{E}}$$

Note that this metric is *real analytic* in its dependence on \mathcal{E} . Also, note that (cf. Proposition 1.4) $V_{\mathcal{E}}$ is naturally isomorphic to the cotangent space to $\mathcal{M}_r(X)$ at the point $[\mathcal{E}]$ of $\mathcal{M}_r(X)$ defined by \mathcal{E} . Thus, we see that we have defined a real analytic Hermitian metric $\langle -, - \rangle_{\mathcal{M}}$ on the cotangent bundle of $\mathcal{M}_r(X)$. We then have the following special case of an important result in modern complex analysis ([Kob], Chapter VII, Theorem 6.36):

Theorem 2.5. *The metric $\langle -, - \rangle_{\mathcal{M}}$ is Kähler. In particular, if we denote by $\mathcal{M}_{[\mathcal{E}]}$ the germ of the complex analytic manifold $\mathcal{M}_r(X)$ at $[\mathcal{E}]$, then we see that this metric defines a canonical local holomorphic embedding*

$$\mathcal{M}_{[\mathcal{E}]} \hookrightarrow V_{\mathcal{E}}^c$$

into the complex conjugate $V_{\mathcal{E}}^c$ of the complex affine space $V_{\mathcal{E}}$.

Proof. The fact that the metric is Kähler is a consequence of [Kob], Chapter VII, Theorem 6.36. (Note that in *loc. cit.*, the quoted result concerns “Einstein-Hermitian metrics.” This condition of “Einstein-Hermitian” simply means in the present context that the *curvature* of the metric is zero. Moreover, by integration, it follows immediately that the curvature is zero if and only if the metric on \mathcal{E} arises from a *constant metric* on $\mathcal{E}|_{\tilde{X}}$, i.e., arises as the natural metric on \mathcal{E}_{ρ} described above.) Finally, whenever one has a real analytic Kähler metric, it follows by general nonsense that one may associate local *canonical coordinates* to the metric (cf. Definition 2.1 of the Introduction to [Mzk1]).
○

This gives us our canonical local coordinates. Moreover, it is easy to see that in the case of $\text{Pic}(X)$, one can define (by taking the integral of an inner product) an analogous Kähler metric on $\text{Pic}(X)$ which gives rise to local holomorphic canonical coordinates that are the same as the coordinates one gets from the exponential sequence. Thus, *the coordinates of Theorem 2.5 are truly the non-abelian analogue of the coordinates on $\text{Pic}(X)$ arising from the exponential exact sequence.*

Remark. Although we are working here with vector bundles with trivial determinant, there is no essential difficulty in generalizing this sort of result to a similar result for G -torsors, where G is an arbitrary reductive algebraic group over \mathbb{C} (cf. the point of view of [Falt3]). Since, however, we have no particular need for such results, we prefer to limit ourselves to the case SL_r .

Finally, just as in Theorem 2.3 of [Mzk1], Introduction, we have the following:

Theorem 2.6. *If we apply $\bar{\partial}$ to the section $s : \mathcal{M}_r(X) \rightarrow \mathcal{A}_r(X)$ of Corollary 2.4, we obtain a $(1,1)$ -form $\bar{\partial}s$ on $\mathcal{M}_r(X)$ which is equal to the $(1,1)$ -form defined by the Kähler metric $\langle -, - \rangle_{\mathcal{M}}$ of Theorem 2.5.*

Proof. The proof is formally the same as that of Theorem 2.3 of the Introduction of [Mzk1]: Indeed, let

$$\rho : \Pi \rightarrow \text{SU}_r$$

be the unitary representation (as in Theorem 2.1) corresponding to some $[\mathcal{E}] \in \mathcal{M}_r(X)$. Then ρ defines a Π -module N , where N is an r -dimensional complex vector space equipped with an Hermitian metric $\langle -, - \rangle_N$ preserved by the action of Π . Let

$$L \stackrel{\text{def}}{=} \text{Ad}(N)$$

denote the trace zero endomorphisms of this vector space N . Thus, L also has a Π -action and a metric preserved by this Π -action. Moreover, L has a *real structure*

$$L_{\mathbf{R}} \subseteq L$$

consisting of those endomorphisms $\epsilon \in N$ that are *anti-Hermitian* with respect to $\langle -, - \rangle_N$ (i.e., if ϵ^* denotes the adjoint of ϵ with respect to this metric, then $\epsilon^* = -\epsilon$). Thus, the group cohomology module $H^1(\Pi, L)$ gets a *natural real structure* $H^1(\Pi, L_{\mathbf{R}})$, hence a natural complex conjugation morphism.

On the other hand, let $\nabla_{\mathcal{E}}$ be the connection on \mathcal{E} of Corollary 2.4. This connection induces a natural connection on $\text{Ad}(\mathcal{E})$. Thus, we may consider the de Rham cohomology module $H_{\text{DR}}^1(X, \text{Ad}(\mathcal{E}))$ of $\text{Ad}(\mathcal{E})$. Moreover, one has a natural *comparison isomorphism* (i.e., the “de Rham isomorphism”)

$$H_{\text{DR}}^1(X, \text{Ad}(\mathcal{E})) \cong H^1(\Pi, L)$$

which thus defines (by the above discussion concerning $H^1(\Pi, L)$) a *complex conjugation morphism*

$$c_{\text{DR}} : H_{\text{DR}}^1(X, \text{Ad}(\mathcal{E})) \rightarrow H_{\text{DR}}^1(X, \text{Ad}(\mathcal{E}))$$

Next, recall (cf. Proposition 1.4) that there is a natural *Hodge filtration* on the de Rham cohomology $H_{\text{DR}}^1(X, \text{Ad}(\mathcal{E}))$: i.e., there is a natural inclusion

$$V_{\mathcal{E}} = H^0(X, \omega_X \otimes_{\mathcal{O}_X} \text{Ad}(\mathcal{E})) \hookrightarrow H_{\text{DR}}^1(X, \text{Ad}(\mathcal{E}))$$

whose cokernel can be identified with the dual of $V_{\mathcal{E}}$. Thus, if one composes this inclusion with c_{DR} and then with the projection to the dual of $V_{\mathcal{E}}$, one obtains a *Hermitian form* on $V_{\mathcal{E}}$. Moreover, just as in Proposition 2.2 of [Mzk1], Introduction, one shows easily that this Hermitian form is equal to the metric $\langle -, - \rangle_{V_{\mathcal{E}}}$ defined above. (Indeed, this follows easily from the definition of the pairing of Serre duality given by integrating forms over X .) On the other hand, note that the real subspace

$$H_{\text{DR}}^1(X, \text{Ad}(\mathcal{E}))^{c_{\text{DR}}=1} \subseteq H_{\text{DR}}^1(X, \text{Ad}(\mathcal{E}))$$

may be naturally identified with the tangent space to the (real analytic submanifold given by the) image of the section s . (Indeed, this follows from the definition given above of the *real structure* $L_{\mathbf{R}}$ on L , and the fact that the Lie algebra of the group of unitary matrices is given by the set of anti-Hermitian matrices.) Thus, it is a *tautology* in linear algebra (cf. the discussion at the beginning of §0.8 of the Introduction to the present work) that $\bar{\partial}s$ is equal to this Hermitian form. Putting everything together, we conclude that $\bar{\partial}s = \langle -, - \rangle_{\mathcal{M}}$, as desired. This completes the proof. \circ

In fact, just as in the case of the moduli of hyperbolic curves, the complex conjugation morphism – read: *Frobenius action* – on the tangent space to $\mathcal{A}_r(X)$ at $\mathrm{Im}(s)$ extends to a Frobenius action $\Phi_{\mathcal{A}}$ on some open neighborhood of $\mathrm{Im}(s)$ in $\mathcal{A}_r(X)$ as follows: Namely, if $(\mathcal{E}, \nabla_{\mathcal{E}})$ defines a point of $\mathcal{A}_r(X)$ corresponding to a representation $\rho : \Pi \rightarrow \mathrm{GL}_r(\mathbf{C})$, then we take

$$\Phi_{\mathcal{A}}(\mathcal{E}, \nabla_{\mathcal{E}})$$

to be the bundle with connection defined by $(\bar{\rho}^t)^{-1}$ (i.e., the inverse of the transposed complex conjugate representation to ρ). Of course, there is the issue of whether or not this new bundle is *stable* – that is why we only claimed that $\Phi_{\mathcal{A}}$ is *defined on some open neighborhood* of $\mathrm{Im}(s)$ in $\mathcal{A}_r(X)$. Moreover, it follows (by considering tangent spaces – cf. the proof of Theorem 2.6 above, as well as the discussion of §0.8 of the Introduction to the present work) that

The fixed point locus of $\Phi_{\mathcal{A}}$ (on some appropriate open neighborhood of $\mathrm{Im}(s)$) is precisely $\mathrm{Im}(s)$.

Finally, it can be shown (cf. Theorem IV.3 of [Falt3]) that the torsor $\mathcal{A}_r(X) \rightarrow \mathcal{M}_r(X)$ is the Hodge-theoretic first Chern class of an ample line bundle on $\mathcal{M}_r(X)$. Thus, we see that we are in the sort of situation discussed in §0.8 of the Introduction to the present work.

§3. The Ordinary p-adic Theory

In this §, we let p be an *odd prime*. We will always assume that the rank r of the vector bundle \mathcal{E} under consideration is *prime to* p .

§3.1. Crystals of Bundles with Connection

Let S be a *noetherian scheme on which p is nilpotent*. Let $X \rightarrow S$ be a *curve*. Then we have the following result:

Proposition 3.1. *There exists a crystal $\Delta_r(X/S)$ on $\text{Crys}(S)$, valued in the category of algebraic spaces, and having the following properties:*

- (1) $\Delta_r(X/S)$ depends (as the notation indicates) only on $X \rightarrow S$ and r ;
- (2) If $(S \hookrightarrow T)$ is an object of $\text{Crys}(S)$ (so the ideal \mathcal{I} defining S in T is equipped with a nilpotent PD-structure – i.e., $\mathcal{I}^{[i]} = 0$ for some i), and $X_T \rightarrow T$ is a curve lifting $X \rightarrow S$, then the value of $\Delta_r(X/S)$ on the object $(S \hookrightarrow T)$ is naturally isomorphic to the algebraic space $\mathcal{A}_r(X_T/T)$ of Theorem 1.3.

Proof. This results immediately from the usual formalism equating crystals (on the one hand) with bundles with connection (on the other). \circ

The formation of the crystal $\Delta_r(X/S)$ on $\text{Crys}(S)$ is also *functorial* in the following sense. Suppose that we are given another curve $Y \rightarrow S$, and an S -morphism

$$\psi : Y \rightarrow X$$

Suppose further that we are given an open subset U of the topological space underlying $\mathcal{A}_r(X/S)$. Since the underlying topological space of a scheme does not change if one passes to a nilpotent thickening, it makes sense to consider the open subspace

$$\Delta_r(X/S)(\tau)|_U \subseteq \Delta_r(X/S)(\tau)$$

for an object $\tau = (S \hookrightarrow T)$ of $\text{Crys}(S)$. The crystal on $\text{Crys}(S)$ (valued in the category of algebraic spaces) thus obtained will be denoted $\Delta_r(X/S; U)$.

Definition 3.2. We shall say that ψ is *U -compatible* if ψ induces a morphism $\psi^* : \mathcal{M}_r(X/S)|_U \rightarrow \mathcal{M}_r(Y/S)$.

Remark. More concretely, the above condition means that if one chooses a stable bundle on X parametrized by U and pulls it back via ψ to a bundle on Y , that bundle will be stable on Y .

Then we have the following result:

Proposition 3.3. *With the above notation, if ψ is U -compatible, then ψ induces a natural morphism $\psi^* : \Delta_r(X/S; U) \rightarrow \Delta_r(Y/S)$ of crystals in algebraic spaces on $\mathrm{Crys}(S)$.*

Proof. This follows by a standard argument (as above) by using the well-known interpretation of vector bundles with connection as crystals on some appropriate crystalline site. \square

§3.2. Frobenius Actions

In this subsection, we assume that S is of characteristic p . As usual, we let $f : X \rightarrow S$ be a curve over S . Let $\Phi_S : S \rightarrow S$ be the absolute Frobenius morphism on S , and

$$\Phi_{X/S} : X \rightarrow X^F \stackrel{\mathrm{def}}{=} X \times_{S, \Phi_S} S$$

(where the structure morphism on the right is Φ_S) the $(S$ -linear) *relative Frobenius* on X . In general, we shall denote the result of pulling back various objects via $(\Phi_S)^i$ by means of a superscript F^i , and we shall denote the i^{th} -power of the relative Frobenius morphism by $\Phi_{X/S}^i$.

Definition 3.4. We shall call a vector bundle \mathcal{E} on X an F -bundle if for some $i \in \mathbf{N}$, there exists an isomorphism

$$\alpha_i : (\Phi_{X/S}^i)^* \mathcal{E}^{F^i} = (\Phi_X^i)^* \mathcal{E} \cong \mathcal{E}$$

We shall say that \mathcal{E} is a *special F -bundle* if \mathcal{E} is special, and the determinant of α_i is the identity.

The theory of “ F -bundles” that we discuss in this subsection will in some sense be a (trivial) special case of the theory of the category \mathcal{MF}^∇ discussed in [Falt1], §2.

One should think of F -bundles as being the parametrized, higher rank analogue of a line bundle of degree zero whose order in the Picard group of X is prime to p . Indeed, it is not difficult to see that when S is the spectrum of a finite field, an F -bundle of rank one is exactly a line bundle of degree zero whose order is prime to p . (More precisely: if the line bundle is annihilated by an integer N , then there exists a positive integer i such that N divides $p^i - 1$; thus, the line bundle is isomorphic to its pull-back by $(\Phi_X^i)^*$.)

Clearly, the property of being an F -bundle is preserved under such basic operations as taking the dual, the tensor product, exterior powers, and symmetric powers. In particular, if \mathcal{E} is an F -bundle, then so is its determinant bundle, so the *relative degree* (with respect to $X \rightarrow S$) of an F -bundle is always zero.

Proposition 3.5. *The restriction of an F -bundle \mathcal{E} to a geometric fiber over S is semistable.*

Proof. Clearly, for the proof, we may assume that S is the spectrum of a finite field. Let $\mathcal{F} \subseteq \mathcal{E}$ be a subbundle. We must show that \mathcal{F} has degree ≤ 0 , so assume $\deg(\mathcal{F}) > 0$. By replacing \mathcal{E} by an appropriate exterior power of itself, we may assume that \mathcal{F} is a *line bundle*. Since pulling back by Φ_X has the effect of raising a line bundle to the power p , we thus obtain that $\mathcal{F}^{\otimes N \cdot p} \subseteq \mathcal{E}$, for any positive integer N . But this implies that

$$\Gamma(X, \mathcal{F}^{\otimes N \cdot p}) \subseteq \Gamma(X, \mathcal{E})$$

which is absurd since $\Gamma(X, \mathcal{E})$ is finite-dimensional over k , while the dimension over k of $\Gamma(X, \mathcal{F}^{\otimes N \cdot p})$ becomes arbitrarily large as $N \rightarrow \infty$. \circ

We remark that Gieseker ([Gie]) has constructed counterexamples to the converse of Proposition 3.5 – cf. also the third Remark following Chapter IV, Theorem 2.3.

One way to obtain F -bundles is as follows. We suppose for convenience that S is *irreducible*. Let

$$Z \rightarrow X$$

be a *finite, étale Galois covering*, where Z is *geometrically connected* over S . This condition that Z be geometrically connected over S may be thought of as expressing the idea that this covering arises from the *relative “geometric” fundamental group* of the curve $X \rightarrow S$, not from some covering of S pulled back to X . Let

$$\Pi \stackrel{\text{def}}{=} \text{Gal}(Z/X)$$

be the *Galois group* of this covering. Let $k \subseteq \Gamma(S, \mathcal{O}_S)$ be a *finite field* contained in $\Gamma(S, \mathcal{O}_S)$. Suppose that we are given a homomorphism

$$\rho : \Pi \rightarrow \text{GL}_r(k)$$

Then Π acts on the vector bundle $\mathcal{E}_Z \stackrel{\text{def}}{=} \mathcal{O}_Z \otimes_k k^n$ by means of the Galois action on Z on the left-hand factor and by means of ρ on the right-hand factor. If we let

$$\mathcal{E}_\rho \stackrel{\text{def}}{=} \{\Pi - \text{invariants of } \mathcal{E}_Z\} \subseteq \mathcal{E}_Z$$

then one sees easily that \mathcal{E}_ρ forms a vector bundle of rank r on X .

Next, observe that we have a commutative diagram

$$\begin{array}{ccccc} Z & \xrightarrow{\Phi_{Z/S}} & Z^F & \longrightarrow & Z \\ \downarrow & & \downarrow & & \downarrow \\ X & \xrightarrow{\Phi_{X/S}} & X^F & \longrightarrow & X \end{array}$$

where the horizontal morphisms on the right are the projections to Z and X , respectively, and both squares are *cartesian*. Note that since all the vertical morphisms are Galois coverings with Galois group Π , we thus obtain an induced *Frobenius automorphism* of Π . Since Π is a *finite set*, some finite power of this Frobenius automorphism will be the *identity*. In other words, the Π -covering $Z \rightarrow X$ (i.e., covering equipped with Π -action) is *necessarily invariant under pull-back* by $(\Phi_{X/S})^i$ (for some i). We may also assume that i is sufficiently large that Φ_S^i is the identity on k . It thus follows from the definition of \mathcal{E}_ρ that $(\Phi_X^i)^*\mathcal{E}_\rho \cong \mathcal{E}_\rho$, i.e., that \mathcal{E}_ρ is an *F-bundle*, hence has geometrically semistable fibres. In what follows, we shall say that “ \mathcal{E}_ρ was constructed from the representation ρ .”

In fact, we can say more:

Proposition 3.6. *With the notations of the previous paragraph, if ρ corresponds to a geometrically irreducible representation of Π on k^r (i.e., the representation remains irreducible after base-change via any morphism $k \rightarrow \Omega$ of k into an algebraically closed field Ω) whose determinant is trivial, then \mathcal{E}_ρ is a special stable vector bundle.*

Proof. Clearly, we may assume that S is the spectrum of a finite field. It then suffices to show that \mathcal{E}_ρ is *stable* (since the construction of \mathcal{E}_ρ commutes with passing to finite extensions of k). Suppose that $\mathcal{F} \subseteq \mathcal{E}_\rho$ is a subbundle of rank m and *degree zero*. Let

$$\mathcal{L} \stackrel{\text{def}}{=} \wedge^m \mathcal{F}; \quad \mathcal{V} \stackrel{\text{def}}{=} \wedge^m \mathcal{E}_\rho$$

Thus, we are given an inclusion $\mathcal{L} \hookrightarrow \mathcal{V}$. Although, a priori, the (necessarily finite – since we are working over a finite field) order of \mathcal{L} may be divisible by p , by pulling back the inclusion $\mathcal{L} \hookrightarrow \mathcal{V}$ several times by Φ_X , we may assume that \mathcal{L} defines a torsion element in $\text{Pic}(X)$ of

order *prime to p*. Also, (after possibly replacing k by a finite extension of k) we may assume that there exists a finite étale Galois covering $U \rightarrow X$, with U geometrically connected over k , such that \mathcal{L} and \mathcal{E}_ρ (and hence also \mathcal{V}) become *trivial* when restricted to U . If Δ is the Galois group of the covering $U \rightarrow X$, then the pulled-back inclusion $\mathcal{L}_U \rightarrow \mathcal{V}_U$ corresponds to some inclusion of Δ -modules

$$\iota : k \hookrightarrow \wedge^m(k^n)$$

Now it is clear from the construction that the point defined by ι in the *projective space of the dual* of $\wedge^m(k^n)$ satisfies the equations that define the Plücker embedding of the Grassmannian of m -dimensional subspaces. (Indeed, this follows from the fact that the inclusion $\mathcal{L} \hookrightarrow \mathcal{V}$ satisfies these equations.) We thus conclude that ι corresponds to a Δ -invariant subspace of k^n of dimension m ; but this contradicts the geometric irreducibility of ρ . \circ

In fact, *essentially every F -bundle can be constructed from a representation of the relative fundamental group*. We make this precise as follows:

Proposition 3.7. *Let \mathcal{E} be an F -bundle on X . Then after replacing S by a finite étale cover of S , there exists a representation ρ as above such that $\mathcal{E} \cong \mathcal{E}_\rho$.*

Proof. Fix an isomorphism

$$\alpha_i : (\Phi_{X/S}^i)^*(\mathcal{E})^{F^i} \cong \mathcal{E}$$

Let K be the field with p^i elements. We may assume that S is a K -scheme of finite type. Let us define a presheaf \mathcal{F} on the étale site of X as follows. If $U \rightarrow X$ is étale, then we let $\mathcal{F}(U)$ be the K -vector space of sections $\sigma \in \Gamma(U, \mathcal{E}|_U)$ such that $\alpha_i(\sigma) = \sigma$. (We shall call such sections *α -invariant*.) It is easy to see that \mathcal{F} is a sheaf in the étale topology. We *claim* that it is, in fact, a *locally constant, constructible sheaf of K -vector spaces of dimension r* . Indeed, it suffices to prove that this is true in an étale neighborhood of every L -rational point $x \in X(L)$, where L is a finite extension of K . But it follows from elementary “semilinear algebra” that (after possibly replacing L by some finite extension of L) the α -invariant subspace of the fibre $\mathcal{E} \otimes_{\mathcal{O}_{X,x}} L$ (of \mathcal{E} at x) is a K -vector space of dimension exactly r . Moreover, it is easy to see that the obstructions to lifting α -invariant sections vanish, and that α -invariant sections lift uniquely, so the α -invariant subspace of the completion $\mathcal{E} \otimes_{\mathcal{O}_{X,x}} \widehat{\mathcal{O}}_{X,x}$ is again a K -vector space of dimension exactly r . By using the fact that “the henselization is henselian,” it follows, moreover, that these α -invariant sections descend uniquely to some étale neighborhood $U \rightarrow X$ of x . This completes the proof of the claim. On the other hand, now

that we know that \mathcal{F} is a local system of K -vector spaces of dimension r , the Proposition follows immediately: Indeed, if we take for the covering “ $Z \rightarrow X$ ” a finite, étale Galois covering that trivializes the local system \mathcal{F} , then the α -invariant subspace of $\Gamma(Z, \mathcal{E}|_Z) \cong \Gamma(S, \mathcal{O}_S)^r$ is isomorphic to K^r and is equipped with a natural action of $\text{Gal}(Z/X)$ (since \mathcal{E} and α_i are obtained by pull-back from X). This thus defines a representation ρ of $\text{Gal}(Z/X)$ into $\text{GL}_r(K)$ such that the resulting \mathcal{E}_ρ is naturally isomorphic to the original vector bundle \mathcal{E} . \circ

We now return to assuming that S is *not necessarily irreducible*. It turns out that F -bundles are easiest to understand when their cohomology acts like the cohomology of an ordinary abelian variety:

Definition 3.8. Suppose that \mathcal{E} is a special, stable F -bundle on X . Then we say that \mathcal{E} is *ordinary* if the morphism

$$\text{Ad}(\alpha_i) : (\Phi_S^i)^* \mathbf{R}^1 f_* \text{Ad}(\mathcal{E}) \rightarrow \mathbf{R}^1 f_* \text{Ad}(\mathcal{E})$$

(induced by α_i) is an isomorphism.

It is easy to see that the condition that $\text{Ad}(\alpha_i)$ be an isomorphism is *independent of the choice of i or α_i* . (Indeed, this follows by considering iterates of $\text{Ad}(\alpha_i)$, and using the fact that since \mathcal{E} is stable, α_i is *unique* up to a constant multiple – cf. the proof of Proposition 1.2.)

To see that ordinary bundles *exist*, we reason as follows. Consider the bundle \mathcal{E}_ρ constructed above in the discussion preceding Proposition 3.6. It is not difficult to see that if the relative Jacobian of Z over S is *ordinary*, then \mathcal{E}_ρ is ordinary. Since it is not difficult to construct such coverings $Z \rightarrow X$ (for instance, by deforming coverings of a totally degenerate stable X that arise “combinatorially,” i.e., from the dual graph of the totally degenerate curve), we thus see that ordinary special, stable F -bundles do indeed exist. This example of an ordinary bundle also further justifies the terminology “ordinary.” We shall study ordinary bundles in greater detail in the following subsection.

§3.3. The Ordinary Case

In this subsection, we study *ordinary stable bundles* on a curve over $W(\overline{\mathbf{F}}_p)$. Thus, for the remainder of this subsection, we fix an odd prime p , a natural number r prime to p , and an artinian local ring A whose residue field k is an *algebraic closure* of \mathbf{F}_p . Let

$$W \stackrel{\text{def}}{=} W(k); \quad V \stackrel{\text{def}}{=} \text{Spec}(W); \quad S \stackrel{\text{def}}{=} \text{Spec}(A);$$

$$S_0 \stackrel{\text{def}}{=} \text{Spec}(k) \subseteq S; \quad k \stackrel{\text{def}}{=} A/\mathfrak{m}$$

Thus, A has a natural structure of W -algebra.

Moreover, we assume that we are given a *curve*

$$f : X \rightarrow V$$

We shall denote the pull-backs of objects over S to objects over S_0 by a subscript zero; more generally, we shall denote pull-backs to W , A , etc. by means of a subscripted W , A , etc. Finally, we fix some base point $x : \text{Spec}(k) \rightarrow X$, and let

$$\Pi \stackrel{\text{def}}{=} \pi_1(X, x)$$

Since S_0 is an \mathbf{F}_p -scheme, we may apply to it the theory of §3.2. In particular, if \mathcal{E} is a special vector bundle of rank r on X_A , it makes sense to say that \mathcal{E}_0 is an ordinary special F -bundle.

Definition 3.9. We shall say that \mathcal{E} is *ordinary* if \mathcal{E}_0 is an ordinary special F -bundle on X_{A_0} .

Let us denote by

$$\text{Bun}_A(X)$$

the category whose objects are *ordinary special stable vector bundles* \mathcal{E} on X_A of rank r , and whose morphisms are morphisms of vector bundles.

Now we consider *continuous representations* of Π . In general, if T is a scheme, we shall denote by $\Gamma(T)$ the topological ring $\Gamma(T, \mathcal{O}_T)$, equipped with the discrete topology.

Definition 3.10. We shall say that ρ is an A -valued representation of Π (of dimension r) if ρ is a continuous homomorphism

$$\rho : \Pi \rightarrow \text{SL}_r(A)$$

Let ρ be an A -valued representation of Π of dimension r . If T is an A -scheme, then we shall denote by ρ_T (respectively, ρ_0) the homomorphism $\Pi \rightarrow \text{SL}_r(\Gamma(T))$ (respectively, $\Pi \rightarrow \text{SL}_r(\Gamma(k))$) obtained by composing with $\text{SL}_r(A) \rightarrow \text{SL}_r(\Gamma(T))$ (respectively, $\text{SL}_r(A) \rightarrow \text{SL}_r(k)$).

We shall say that ρ is *irreducible* if ρ_0 is irreducible (cf. Proposition 3.6). We shall say that ρ is *ordinary* if the first continuous group cohomology module (relative to the action of Π on k^r obtained from ρ_0) satisfies:

$$\dim_k \{H^1(\Pi, \text{Ad}(k^r))\} \geq (r^2 - 1)(g - 1)$$

Let us denote by

$$\text{Rep}_A(X)$$

the category of *ordinary, irreducible* A -valued representations of Π of dimension r . Now observe that we have a functor

$$\Psi_A : \text{Rep}_A(X) \rightarrow \text{Bun}_A(X)$$

given by constructing a stable bundle out of a representation of Π just as in the discussion preceding Proposition 3.6. Indeed, to see that this construction yields bundles which do indeed constitute objects of $\text{Bun}_A(X)$, it suffices to apply Proposition 3.6, and then to show that the resulting bundle is *ordinary*. To show that the resulting bundle is ordinary, we reason as follows: First, observe that by thinking of “ H^1 ” in terms of *extensions* and applying Proposition 3.7, it follows that the cohomology module

$$H^1(\Pi, \text{Ad}(k^r))$$

may be identified with the k -subspace of $H^1(X, \text{Ad}(\mathcal{E}))$ on which all iterates of $\text{Ad}(\alpha_i)$ act *bijectively* (cf. Definition 3.8). Thus, we see that the condition that the *representation be ordinary* corresponds precisely to the condition that *the bundle constructed out of the representation be ordinary*. Moreover, it follows easily from the definitions (and pulling back to a finite étale Galois covering of X_A on which the bundles involved are trivial) that Ψ_A is *fully faithful*.

Suppose that \mathcal{E} is a bundle on X_A which is in the *essential image* of Ψ_A . Then it follows that there exists a finite, connected, étale Galois covering

$$Z \rightarrow X$$

(with Galois group $\Gamma \stackrel{\text{def}}{=} \text{Gal}(Z/X)$) such that $\mathcal{E}|_{Z_A}$ becomes *trivial*. Moreover, since Γ acts on $\mathcal{E}|_{Z_A}$ by matrices which are *constant* (relative to the morphism $X_A \rightarrow S$), it follows that the the trivial connection on $\mathcal{E}|_{Z_A}$ is fixed by the action of Γ , hence descends to a *canonical connection* $\nabla_{\mathcal{E}}$ on \mathcal{E} . Next, let us note that

$$(\mathcal{E}, \nabla_{\mathcal{E}})$$

defines a *crystal* on $\text{Crys}(((X_A)_{\mathbf{F}_p})/S)$. On the other hand, since the reduction $X_{\mathbf{F}_p}$ of X modulo p is defined over some *finite field*, it follows that for some large i , $X_{\mathbf{F}_p}^{F^i} = X_{\mathbf{F}_p}$. Thus, for some large i , we have a k -linear Frobenius morphism $\Phi_X^i : X_{\mathbf{F}_p} \rightarrow X_{\mathbf{F}_p}$, which we can base change via $W \rightarrow A$ to obtain an S -linear morphism

$$(\Phi_X^i)_A : (X_A)_{\mathbf{F}_p} \rightarrow (X_A)_{\mathbf{F}_p}$$

Thus, we can speak of $(\Phi_X^i)_A^*(\mathcal{E}, \nabla_{\mathcal{E}})$ (i.e., the pull-back of the crystal $(\mathcal{E}, \nabla_{\mathcal{E}})$ by $(\Phi_X^i)_A$).

Lemma 3.11. *For some sufficiently large i , we have $(\Phi_X^i)_A^*(\mathcal{E}, \nabla_{\mathcal{E}}) \cong (\mathcal{E}, \nabla_{\mathcal{E}})$.*

Proof. Choose i large enough so that $Z_{\mathbf{F}_p}^{F^i} = Z_{\mathbf{F}_p}$. Thus, we have a k -linear Frobenius morphism $\Phi_Z^i : Z_{\mathbf{F}_p} \rightarrow Z_{\mathbf{F}_p}$, which covers $\Phi_X^i : X_{\mathbf{F}_p} \rightarrow X_{\mathbf{F}_p}$. This gives us an action of the i^{th} power of Frobenius on the quotient $\Pi \rightarrow \Gamma$. Since Γ is finite, by choosing i sufficiently large, we may even assume that the i^{th} power of Frobenius acts trivially on Γ . Since $(\Phi_X^i)_A$ is A -linear, it thus follows trivially from the *definition of the functor* Ψ_A that $(\Phi_X^i)_A^*(\mathcal{E}, \nabla_{\mathcal{E}}) \cong (\mathcal{E}, \nabla_{\mathcal{E}})$. \circ

Now let us consider the *de Rham cohomology* of $\text{Ad}(\mathcal{E}_0, \nabla_{\mathcal{E}_0})$. First, note that we have an exact sequence

$$\begin{aligned} 0 \rightarrow H^0(X_{\mathbf{F}_p}, \text{Ad}(\mathcal{E}_0) \otimes \omega_{X_{\mathbf{F}_p}/k}) &\rightarrow H_{\text{DR}}^1(X_{\mathbf{F}_p}, \text{Ad}(\mathcal{E}_0, \nabla_{\mathcal{E}_0})) \\ &\rightarrow H^1(X_{\mathbf{F}_p}, \text{Ad}(\mathcal{E}_0)) \rightarrow 0 \end{aligned}$$

(arising from the Hodge filtration – cf. Proposition 1.4). Let us write

$$H_{\text{DR}} \stackrel{\text{def}}{=} H_{\text{DR}}^1(X_{\mathbf{F}_p}, \text{Ad}(\mathcal{E}_0, \nabla_{\mathcal{E}_0}))$$

Then let us observe that the fact that $(\mathcal{E}_0, \nabla_{\mathcal{E}_0})$ is preserved by some power, say the i^{th} power, of Frobenius (cf. Lemma 3.11) means that there is a naturally induced action

$$\Phi_H^i : (H_{\text{DR}})^{F^i} \rightarrow H_{\text{DR}}$$

of the i^{th} -power of Frobenius on H_{DR} . Note that, relative to the Hodge filtration reviewed above, $\Phi_H^i(F^1(H_{\text{DR}})^{F^i}) = 0$. (Indeed, this follows from the fact that $F^1(-)$ is made up of *differentials* – cf. the above exact sequence – and the fact that differentials in characteristic p vanish under pull-back by Frobenius.) Moreover, by ordinariness (cf. the

above exact sequence, and Definition 3.8), Φ_H^i maps F^0/F^1 isomorphically onto F^0/F^1 .

We are now ready to state and prove the following result (which is the p -adic analogue of Theorem 2.1 and its corollaries in the complex case):

Theorem 3.12. *The functor*

$$\Psi_A : \operatorname{Rep}_A(X) \rightarrow \operatorname{Bun}_A(X)$$

is an equivalence of categories. In particular, every object \mathcal{E} of $\operatorname{Bun}_A(X)$ admits a unique (canonical) connection characterized by the property that the crystal on $\operatorname{Crys}(((X_A)_{\mathbf{F}_p})/A)$ defined by $(\mathcal{E}, \nabla_{\mathcal{E}})$ is fixed by some power of Frobenius.

Proof. We have already observed above that Ψ_A is *fully faithful*. Let us prove that Ψ_A is *essentially surjective*. We use induction on the length of A . When $A = k$, the result follows immediately from Proposition 3.7. Now we proceed to the induction step. Thus, we consider a *surjection*

$$A \rightarrow B$$

whose kernel is an ideal of A of length one. We will denote the result of base-changing objects over A to objects over B by means of a subscript B .

Let \mathcal{E}_A be an *object* of $\operatorname{Bun}_A(X)$. By the induction hypothesis, it follows that \mathcal{E}_B is in the *essential image* of Ψ_B . Let us denote by \mathcal{A} the étale local system of k -vector spaces on $X_{\mathbf{F}_p}$ defined by the Π -module $\operatorname{Ad}(k^r)$. Note that since the action of Π is continuous, it follows that for some finite field K (with, say, p^i elements), the action of Π on k^r arises from an action on K^r . Thus, we obtain a local system \mathcal{A}_K such that $\mathcal{A}_K \otimes_K k = \mathcal{A}$. Let

$$\mathcal{A}_X \stackrel{\text{def}}{=} \mathcal{A}_K \otimes_K \mathcal{O}_{X_{\mathbf{F}_p}}$$

Then note that \mathcal{A}_X is equipped with a natural action $\Phi_{\mathcal{A}}^i$ of the i^{th} power of Frobenius such that the sheaf of $\Phi_{\mathcal{A}}^i$ -invariant sections of \mathcal{A}_X is given by \mathcal{A}_K (cf. the proof of Proposition 3.7). Thus, the (Artin-Schreier type) short exact sequence of sheaves

$$0 \longrightarrow \mathcal{A}_K \longrightarrow \mathcal{A}_X \xrightarrow{1-\Phi_{\mathcal{A}}^i} \mathcal{A}_X \longrightarrow 0$$

on the étale site of $X_{\mathbf{F}_p}$ gives rise to a long exact sequence of étale cohomology groups. This long exact sequence then defines an exact sequence:

$$\begin{aligned} \text{Coker}\{(1 - \Phi_{\mathcal{A}}^i) : H_{\text{et}}^1(X_{\mathbf{F}_p}, \mathcal{A}_X) \rightarrow H_{\text{et}}^1(X_{\mathbf{F}_p}, \mathcal{A}_X)\} \\ \rightarrow H_{\text{et}}^2(X_{\mathbf{F}_p}, \mathcal{A}_K) \rightarrow H_{\text{et}}^2(X_{\mathbf{F}_p}, \mathcal{A}_X) \end{aligned}$$

On the other hand, since \mathcal{A}_X is a *coherent* \mathcal{O}_X -module, we have

$H_{\text{et}}^2(X_{\mathbf{F}_p}, \mathcal{A}_X) = 0$. Moreover, by the *ordinariness condition* (cf. Definition 3.8, and the fact that \mathcal{A}_X may be identified with $\text{Ad}(\mathcal{E}_0)$), it follows (by elementary “semilinear algebra”) that the endomorphism $1 - \Phi_{\mathcal{A}}^i$ of $H_{\text{et}}^1(X_{\mathbf{F}_p}, \mathcal{A}_X)$ is *surjective*. Thus, we conclude that

$$H_{\text{et}}^2(X_{\mathbf{F}_p}, \mathcal{A}_K) = 0$$

But this implies (by the general nonsense of deformation theory) that the representation of Π into $\text{SL}_r(B)$ that defines \mathcal{E}_B *lifts* to a representation of Π into $\text{SL}_r(A)$. In particular, it follows that (if $\nabla_{\mathcal{E}_B}$ is the canonical connection $\nabla_{\mathcal{E}_B}$ on \mathcal{E}_B – cf. the discussion preceding Lemma 3.11) then $(\mathcal{E}_B, \nabla_{\mathcal{E}_B})$ *lifts to some A -flat $(\mathcal{F}, \nabla_{\mathcal{F}})$* .

Now let us consider the *difference* between the two deformations \mathcal{F} and \mathcal{E}_A (on X_A) of the bundle \mathcal{E}_B on X_B . This difference defines an element

$$\delta \in H^1(X_{\mathbf{F}_p}, \text{Ad}(\mathcal{E}_0)) = (F^0/F^1)(H_{\text{DR}})$$

Note that $H^1(X_{\mathbf{F}_p}, \text{Ad}(\mathcal{E}_0))$ is a k -vector space. On the other hand, since $X_{\mathbf{F}_p}$ and \mathcal{E}_0 are, in fact, defined over some finite field, it thus follows that δ is also defined over some finite field, hence is *fixed by the action of* Φ_H^i , for some i . Thus, by lifting δ to an element of H_{DR} that is fixed by Φ_H^i , and adding (the deformation represented by) this element to $(\mathcal{F}, \nabla_{\mathcal{F}})$, we see that there exists a unique connection $\nabla_{\mathcal{E}_A}$ on \mathcal{E}_A such that $(\mathcal{E}_A, \nabla_{\mathcal{E}_A})$ is fixed by the i^{th} power of Frobenius. Then by taking i^{th} -power Frobenius invariants of $(\mathcal{E}_A, \nabla_{\mathcal{E}_A})$ (cf. the proof of Proposition 3.7), it follows that \mathcal{E}_A lies in the essential image of Ψ_A . This completes the proof. \square

§3.4. Canonical Coordinates via the Weil Conjectures

In this subsection, we let k be the *finite field* of $q = p^f$ elements. Let

$$W \stackrel{\text{def}}{=} W(k); \quad V \stackrel{\text{def}}{=} \text{Spec}(W)$$

and let

$$X \rightarrow V$$

be a *curve*. Let

$$(\mathcal{E}, \nabla_{\mathcal{E}})$$

be an *ordinary, special stable vector bundle with connection* of rank r on X . Let us assume, moreover, that $(\mathcal{E}, \nabla_{\mathcal{E}})$ becomes *trivial* upon restriction to a finite étale covering of X which is geometrically connected over W . Note that $(\mathcal{E}, \nabla_{\mathcal{E}})$ defines a *crystal* on $\mathrm{Crys}(X_{\mathbf{F}_p}/W)$. Let us assume that

$$(\Phi_X^f)^*(\mathcal{E}, \nabla_{\mathcal{E}}) \cong (\mathcal{E}, \nabla_{\mathcal{E}})$$

Let us denote by

$$\mathcal{A}_{\mathcal{E}} \rightarrow V \quad (\text{respectively, } \mathcal{M}_{\mathcal{E}} \rightarrow V)$$

the *completion* of $\mathcal{A}_r(X/V)$ (respectively, $\mathcal{M}_r(X/V)$) – as in Theorem 1.3 – at the V -valued point defined by $(\mathcal{E}, \nabla_{\mathcal{E}})$ (respectively, \mathcal{E}). Thus, $\mathcal{A}_{\mathcal{E}}$ (respectively, $\mathcal{M}_{\mathcal{E}}$) is formally smooth over W of relative dimension $2(r^2 - 1)(g - 1)$ (respectively, $(r^2 - 1)(g - 1)$).

Since $(\Phi_X^f)^*$ *preserves* $(\mathcal{E}, \nabla_{\mathcal{E}})$, it follows from Propositions 3.1 and 3.3 that $(\Phi_X^f)^*$ induces a *Frobenius action*

$$\Phi_{\mathcal{A}}^f : \mathcal{A}_{\mathcal{E}} \rightarrow \mathcal{A}_{\mathcal{E}}$$

on the formal scheme $\mathcal{A}_{\mathcal{E}}$. Let K be the quotient field of W . Let us write

$$(\mathcal{A}_{\mathcal{E}})_K; \quad (\mathcal{M}_{\mathcal{E}})_K$$

for the completions of $\mathcal{A}_r(X/V) \otimes_W K$ and $\mathcal{M}_r(X/V) \otimes_W K$, respectively, at the K -valued points defined by $(\mathcal{E}, \nabla_{\mathcal{E}}) \otimes_W K$ and $\mathcal{E} \otimes_W K$. Thus, $\Phi_{\mathcal{A}}^f$ also induces a Frobenius action on $(\mathcal{A}_{\mathcal{E}})_K$, which we denote by $\Phi_{\mathcal{A}_K}^f$.

Next, let us note that as a consequence of Theorem 3.12, we have the following (cf. Corollary 2.4 in the complex case):

Corollary 3.13. *There is a unique section*

$$s : \mathcal{M}_{\mathcal{E}} \rightarrow \mathcal{A}_{\mathcal{E}}$$

of the natural morphism $\mathcal{A}_{\mathcal{E}} \rightarrow \mathcal{M}_{\mathcal{E}}$ whose image is stabilized by $\Phi_{\mathcal{A}}^f$.

Proof. The section $s : \mathcal{M}_{\mathcal{E}} \rightarrow \mathcal{A}_{\mathcal{E}}$ is defined by assigning to a deformation \mathcal{F} of \mathcal{E} the pair $(\mathcal{F}, \nabla_{\mathcal{F}})$, where $\nabla_{\mathcal{F}}$ is the *canonical connection* of Theorem 3.12. Note that although Theorem 3.12 concerned the case where the base A was artinian with residue field equal to $\overline{\mathbf{F}}_p$, it is easy to see that one may extend Theorem 3.12 immediately to bases like $\mathcal{M}_{\mathcal{E}}$ by:

- (1) passing to inverse limits, and
- (2) observing that since the canonical connection is *canonical* (i.e., uniquely characterized by the property of being invariant under a power of Frobenius), it is invariant under the natural action of $\text{Gal}(\overline{\mathbf{F}}_p/k)$, hence *descends* back down to $\mathcal{M}_{\mathcal{E}}$ from $\mathcal{M}_{\mathcal{E}} \otimes_{W(k)} W(\overline{\mathbf{F}}_p)$ (topological tensor product).

To see that $\text{Im}(s)$ is *stabilized* by Frobenius, it suffices to observe that the Frobenius pull-back of a crystal defined by some canonical pair $(\mathcal{F}, \nabla_{\mathcal{F}})$ (i.e., a pair for which $\nabla_{\mathcal{F}}$ is canonical) is itself always canonical. Indeed, this follows from the definitions (i.e., that the “canonical connection” is the *unique* – cf. Theorem 3.12 – connection fixed by a power of Frobenius) and the fact that any power of Frobenius commutes with Frobenius itself! \circ

It thus follows that $\Phi_{\mathcal{A}}^f$ induces a *Frobenius action*

$$\Phi_{\mathcal{M}}^f : \mathcal{M}_{\mathcal{E}} \rightarrow \mathcal{M}_{\mathcal{E}}$$

(such that the section s is Frobenius-equivariant). Similarly, we have an action of Frobenius on $(\mathcal{M}_{\mathcal{E}})_K$. Naturally, $\Phi_{\mathcal{M}}^f$ fixes the point $[\mathcal{E}] \in \mathcal{M}_{\mathcal{E}}(W)$ defined by \mathcal{E} . Let

$$\Omega_{\mathcal{E}}$$

be the *cotangent space* to $\mathcal{M}_{\mathcal{E}}$ at $[\mathcal{E}]$. Thus, $\Omega_{\mathcal{E}}$ is dual to $H^1(X, \text{Ad}(\mathcal{E}))$, hence is a free W -module of rank $(r^2 - 1)(g - 1)$. Moreover, $\Phi_{\mathcal{M}}^f$ induces a (W -linear) *Frobenius action*

$$\Phi_{\Omega_{\mathcal{E}}}^f : \Omega_{\mathcal{E}} \rightarrow \Omega_{\mathcal{E}}$$

whose reduction modulo p is none other than the dual to the morphism appearing in Definition 3.8 (for $i = f$). Thus, it follows that $\Phi_{\Omega_{\mathcal{E}}}^f$ is an *isomorphism*. In particular, we have the following result:

Proposition 3.14. *The morphism $\Phi_{\mathcal{M}}^f : \mathcal{M}_{\mathcal{E}} \rightarrow \mathcal{M}_{\mathcal{E}}$ is an isomorphism.*

Next, we would like to consider the *eigenvalues of the Frobenius action on $\Omega_{\mathcal{E}}$* . To do this, first note that

$$\Omega_{(\mathcal{E}, \nabla_{\mathcal{E}})} = H_{\text{DR}}^1(X, \text{Ad}(\mathcal{E}, \nabla_{\mathcal{E}}))$$

(i.e., the cotangent space to $\mathcal{A}_{\mathcal{E}}$ at the W -valued point $[(\mathcal{E}, \nabla_{\mathcal{E}})]$) also gets a natural Frobenius action, induced by $\Phi_{\mathcal{A}}^f$. Moreover, let us note that (by differentiating the section s of Corollary 3.13, we see that) $\Omega_{\mathcal{E}}$ appears *Frobenius-equivariantly* as a *quotient* of the W -module $H_{\text{DR}}^1(X, \text{Ad}(\mathcal{E}, \nabla_{\mathcal{E}}))$ (where $\nabla_{\mathcal{E}}$ is the canonical connection).

Let $Y \rightarrow X$ be a finite étale covering which is geometrically connected over W such that $(\mathcal{E}, \nabla_{\mathcal{E}})|_Y$ is *trivial*. (Note that such a $Y \rightarrow X$ exists by the assumptions made at the beginning of this subsection.) Then we have a *Frobenius-equivariant W -linear inclusion*

$$H_{\text{DR}}^1(X, \text{Ad}(\mathcal{E}, \nabla_{\mathcal{E}})) \hookrightarrow H_{\text{DR}}^1(Y, \mathcal{O}_Y)^{r^2-1}$$

(where the superscripted $r^2 - 1$ denotes “the direct sum of $r^2 - 1$ copies of”). Let us denote the (W -linear) action of the f^{th} power of Frobenius on $H_{\text{DR}}^1(Y, \mathcal{O}_Y)$ by Φ_Y^f . Then one knows by the *Riemann hypothesis* part of the Weil Conjectures for curves (see, e.g., [AV] for a proof) that the *eigenvalues of Φ_Y^f are algebraic integers whose complex absolute value at every infinite place is always equal to $q^{\frac{1}{2}}$* . It follows in particular that the action of Φ_Y^f on $H_{\text{DR}}^1(Y, \mathcal{O}_Y)$ has *no eigenvalues in common with the action of $\mathbf{S}^N(\Phi_Y^f)$ on $\mathbf{S}^N(H_{\text{DR}}^1(Y, \mathcal{O}_Y))$, for any $N \geq 2$* . Thus, we obtain the following:

Lemma 3.15. *The action of $\Phi_{\Omega_{\mathcal{E}}}^f$ on $\Omega_{\mathcal{E}}$ has no eigenvalues in common with the action of $\mathbf{S}^N(\Phi_{\Omega_{\mathcal{E}}}^f)$ on $\mathbf{S}^N(\Omega_{\mathcal{E}})$, for any $N \geq 2$.*

We are now ready to construct *canonical coordinates on $\mathcal{M}_{\mathcal{E}}$* . First, let us write

$$(\mathcal{M}_{\mathcal{E}})_K = \text{Spf}(R_K)$$

The point $[\mathcal{E}]$ defines an *augmentation* $R_K \rightarrow K$ whose kernel we denote by

$$I_K \subseteq R_K$$

Write Ψ_R for the action induced by $\Phi_{\mathcal{M}}^f$ on R_K . Thus, Ψ_R is K -linear and preserves I_K . Moreover, we have a *Frobenius-equivariant isomorphism*

$$I_K/I_K^2 \cong (\Omega_{\mathcal{E}})_K$$

Thus, since I_K^2/I_K^N is (Frobenius-equivariantly) a successive extension of various symmetric powers $\mathbf{S}^{N'}(\Omega_{\mathcal{E}})_K$, for $2 \leq N' \leq N$, it follows from Lemma 3.15 (and elementary linear algebra) that there is a unique Frobenius-equivariant section $(\Omega_{\mathcal{E}})_K \rightarrow I_K/I_K^N$ of the surjection

$$I_K/I_K^N \rightarrow I_K/I_K^2 = (\Omega_{\mathcal{E}})_K$$

(for all $N \geq 2$). Since I_K is the *inverse limit* (over N) of the I_K/I_K^N 's, it thus follows that we obtain the following result (analogous to Theorem 2.5 in the complex case – cf. also the discussion of Chapter VIII, §2.4):

Theorem 3.16. *There exists a unique Frobenius-equivariant K -linear section*

$$\kappa : (\Omega_{\mathcal{E}})_K \rightarrow I_K$$

of $I_K \rightarrow I_K/I_K^2 = (\Omega_{\mathcal{E}})_K$. Thus, the functions in the image of κ are canonical coordinates on $\mathcal{M}_{\mathcal{E}}$.

Moreover, just as in the complex case, *these canonical coordinates generalize the canonical coordinates on $\mathrm{Pic}(X)$ obtained by means of the p -adic exponential map at the identity of $\mathrm{Pic}(X)$.*

Bibliography

- [Abik] W. Abikoff, *The Real Analytic Theory of Teichmüller Space*, Lecture Notes in Mathematics **820**, Springer (1980).
- [Atiy] M. F. Atiyah, Commentary on the Article of Manin, in *Arbeitstagung Bonn 1984*, Lecture Notes in Mathematics **1111**, Springer (1985).
- [AV] D. Mumford, *Abelian Varieties*, Oxford Univ. Press (1974).
- [Bers] L. Bers, Quasiconformal Mappings and Teichmüller's Theorem, in *Analytic Functions*, Princeton Univ. Press (1960), pp. 89-119.
- [BHC] A. Borel and Harish-Chandra, Arithmetic Subgroups of Algebraic Groups, *Ann. Math.* **75**, No. 2, pp. 485-535 (1962).
- [BK] S. Bloch and K. Kato, L-Functions and Tamagawa Numbers, in *The Grothendieck Festschrift*, Volume I, Birkhäuser (1990), pp. 333-400.
- [BO] P. Berthelot and A. Ogus, *Notes on Crystalline Cohomology*, Princeton Univ. Press (1978).
- [Brooks] R. Brooks, Circle Packings and Co-compact Extensions of Kleinian Groups, *Invent. Math.* **86**, pp. 461-469 (1986).
- [Cara] H. Carayol, Sur la mauvaise réduction des courbes de Shimura, *Compositio Math.* **59**, No. 2, pp. 151-230 (1986).
- [CF] J. W. S. Cassels and A. Froehlich, *Algebraic Number Theory*, Academic Press (1967).
- [DM] P. Deligne and D. Mumford, The Irreducibility of the Moduli Space of Curves of Given Genus, *IHES Publ. Math.* **36**, pp. 75-109 (1969).
- [EV] H. Esnault and E. Viehweg, Effective Bounds for Semi-positive Sheaves and for the Height of Points of Curves over Complex Function Fields, *Compositio Math.* **76**, pp. 69-85 (1990).

- [Falt1] G. Faltings, Crystalline Cohomology and p -adic Galois Representations, *Proceedings of the First JAMI Conference*, Johns Hopkins Univ. Press (1990), pp. 25-79.
- [Falt2] G. Faltings, Real Projective Structures on Riemann Surfaces, *Compositio Math.* **48**, Fasc. 2, pp. 223-269 (1983).
- [Falt3] G. Faltings, Stable G -Bundles and Projective Connections, *Journal of Algebraic Geometry* **2**, No. 3, pp. 507-568 (1993).
- [FC] G. Faltings and C.-L. Chai, *Degenerations of Abelian Varieties*, Springer (1990).
- [FM] W. Fulton and D. Mumford, On the Irreducibility of the Moduli Space of Curves, *Invent. Math.* **67**, pp. 87-88 (1982).
- [Font1] J.-M. Fontaine, Cohomologie de de Rham, cohomologie cristalline et représentations p -adiques, *Algebraic Geometry, Proceedings, Tokyo/Kyoto 1982*, Lecture Notes in Mathematics **1016**, Springer (1983), pp. 86-108.
- [Font2] J.-M. Fontaine, *Groupes p -divisibles sur les corps locaux*, *Astérisque* **47-48**, Société Mathématique de France (1977).
- [Fulton1] W. Fulton, *Intersection Theory*, *Ergebnisse der Mathematik und ihrer Grenzgebiete*, 3. Folge, Band **2**, Springer (1984).
- [Fulton2] W. Fulton, Hurwitz Schemes and the Irreducibility of Moduli of Algebraic Curves, *Ann. Math.* **90**, pp. 542-575 (1969).
- [Gard] F. Gardiner, *Teichmüller Theory and Quadratic Differentials*, John Wiley and Sons (1987).
- [Geer] G. van der Geer, *Cycles on the Moduli Space of Abelian Varieties*, preprint of the Department of Mathematics, Kyoto Univ. 96-10 (June 1996).
- [Gie] D. Gieseker, Stable Vector Bundles and the Frobenius Morphism, *Ann. Sci. École Norm. Sup.* **6**, pp. 95-101 (1973).
- [Gunning] R. C. Gunning, Special Coordinate Coverings of Riemann Surfaces, *Math. Annalen* **170**, pp. 67-86 (1967).
- [Harts1] R. Hartshorne, *Ample Subvarieties of Algebraic Varieties*, Lecture Notes in Mathematics **156**, Springer (1970).
- [Harts2] R. Hartshorne, *Algebraic Geometry*, Graduate Texts in Mathematics **52**, Springer (1977).
- [HZ] J. Harer and D. Zagier, The Euler Characteristic of the Moduli Space of Curves, *Invent. Math.* **67**, pp. 457-485 (1982).

- [Ih1] Y. Ihara, Schwarzian Equations, *Jour. Fac. Sci. Univ. Tokyo*, Sect IA Math. **21**, pp. 97-118 (1974).
- [Ih2] Y. Ihara, On the Differentials Associated to Congruence Relations and the Schwarzian Equations Defining Uniformizations, *Jour. Fac. Sci. Univ. Tokyo*, Sect. IA Math. **21**, pp. 309-332 (1974).
- [Ih3] Y. Ihara, On the Frobenius Correspondences of Algebraic Curves, in *Algebraic Number Theory*, papers contributed for the International Symposium, Kyoto, 1976, Japan Soc. Prom. Sci. (1977).
- [Ih4] Y. Ihara, Lifting Curves over Finite Fields Together with the Characteristic Correspondence $\Pi + \Pi'$, *Jour. of Algebra* **75**, No. 2, pp. 452-483 (1982).
- [Kato] K. Kato, Logarithmic Structures of Fontaine-Illusie, *Proceedings of the First JAMI Conference*, Johns Hopkins Univ. Press (1990), pp. 191-224.
- [Katz] N. Katz, p-adic Interpolation of Real Analytic Eisenstein Series, *Ann. Math.* **104**, pp. 459-571 (1976).
- [KM] N. Katz and B. Mazur, *Arithmetic Moduli of Elliptic Curves*, Annals of Mathematics Studies **108**, Princeton Univ. Press (1985).
- [Knud] F. F. Knudsen, The Projectivity of the Moduli Space of Stable Curves, III, *Math. Scand.* **52**, pp. 200-212 (1983).
- [Kob] S. Kobayashi, *The Differential Geometry of Complex Vector Bundles*, Publications of the Mathematical Society of Japan **15**, Iwanami Shoten Publishers and Princeton Univ. Press (1987).
- [Lehto] O. Lehto, *Univalent Functions and Teichmüller Spaces*, Graduate Texts in Mathematics **109**, Springer (1987).
- [Manin] Yu. Manin, New Dimensions in Geometry, in *Arbeitstagung Bonn 1984*, Lecture Notes in Mathematics **1111**, Springer (1985).
- [Mask] B. Maskit, *Kleinian Groups*, Grundlehren der mathematischen Wissenschaften **287**, Springer (1988).
- [Mats] H. Matsumura, *Commutative Algebra (Second Edition)*, The Benjamin/Cummings Publishing Company (1980).
- [Mess] W. Messing, *The Crystals Associated to Barsotti-Tate Groups; with Applications to Abelian Schemes*, Lecture Notes in Mathematics **264**, Springer (1972).
- [Mumf1] D. Mumford, Stability of Projective Varieties, *Enseignement Math.* **23**, pp. 39-110 (1977).

- [Mumf2] D. Mumford, An Analytic Construction of Degenerating Curves over Complete Local Rings, *Compositio Math.* **24**, pp. 129-174 (1972).
- [Mzk1] S. Mochizuki, A Theory of Ordinary p -adic Curves, *Publ. RIMS, Kyoto Univ.* **32**, No. 6, pp. 957-1151 (1996).
- [Mzk2] S. Mochizuki, The Geometry of the Compactification of the Hurwitz Scheme, *Publ. RIMS, Kyoto Univ.* **31**, No. 3, pp. 355-441 (1995).
- [Mzk3] S. Mochizuki, *The Local Pro- p Anabelian Geometry of Curves*, RIMS Preprint 1097 (August 1996), 84 pp; to appear in *Invent. Math.*
- [Mzk4] S. Mochizuki, Correspondences on Hyperbolic Curves, *Jour. Pure Appl. Algebra* **131**, No. 3, pp. 227-244 (1998).
- [Mzk5] S. Mochizuki, *The Generalized Ordinary Moduli of p -adic Hyperbolic Curves*, RIMS Preprint 1051 (December 1995); 281 pp.
- [Mzk6] S. Mochizuki, *Combinatorialization of p -adic Teichmüller Theory*, RIMS Preprint 1076 (April 1996); 32 pp.
- [Naka] H. Nakamura, Galois Rigidity of Profinite Fundamental Groups, *Sugaku Expositions* **10**, No. 2, pp. 195-215 (1997).
- [Oda] T. Oda, *Étale Homotopy Type of the Moduli Spaces of Algebraic Curves*, preprint (1990).
- [Oort] F. Oort, Complete Subvarieties of Moduli Spaces, in *Abelian Varieties* (W. Barth, K. Hulek, H. Lange, eds.), de Gruyter Verlag, Berlin (1995), pp. 225-235.
- [Serre] J.-P. Serre, *Local Fields*, Graduate Texts in Mathematics **67**, Springer (1979).
- [Sesh] C. S. Seshadri, *Fibrés vectoriels sur les courbes algébriques*, *Astérisque* **96**, Société Mathématique de France (1982).
- [Shi] G. Shimura, *Introduction to the Arithmetic Theory of Automorphic Forms*, *Publ. Math. Soc. of Japan* **11**, Iwanami Shoten and Princeton Univ. Press (1971).
- [Surf] D. Mumford, *Lectures on Curves on an Algebraic Surface*, *Annals of Mathematics Studies* **59**, Princeton Univ. Press (1966).
- [Szp] L. Szpiro, *Séminaire sur les pinceaux de courbes de genre au moins deux*, *Astérisque* **86**, Société Mathématique de France (1981).
- [Take] K. Takeuchi, Arithmetic Fuchsian Groups with Signature $(1; e)$, *Jour. Math. Soc. Japan* **35**, No. 3, pp. 381-407 (1983).

- [Thur] W. P. Thurston, Three Dimensional Manifolds, Kleinian Groups and Hyperbolic Geometry, *Bull of the Amer. Math. Soc.* **6**, No. 3, pp. 357-381 (1982).
- [Vojta] P. Vojta, *Diophantine Approximations and Value Distribution Theory*, Lecture Notes in Mathematics **1239**, Springer (1987).

Index

- active (bundle) — II, §1.1
- active for Π — III, §1.1
- active node — V, §1
- adjustable, -ment — II, §1.6
- admissible — Intro, §1.1; II, §1.1
- affine stack — III, §2.3
- affinization — IX, §1.3
- anabelian geometry — Intro, §1.6; VIII, §2.6
- anabelian VF-pattern — IV, §3.1
- aphilial (bundle) — V, §2.1
- aphilial node — Intro, §1.2; V, §1
- arithmetic Frobenius (Kähler) venues — Intro, §0.8
- arithmetic Kodaira-Spencer problem — Intro, §2.3
- arithmetic Schwarz torsor — Intro, §2.3.2
- atom — Intro, §1.2; V, §0
- auxiliary bundle — VI, §1.4

- Bers coordinates — Intro, §0.2
- binary ordinary Frobenius system — VIII, §2.1
- binary ordinary geometry — Intro, §1.6
- binary VF-pattern — Intro, §1.4; III, §1.1

- canonical affine parameters — VIII, §2.4
- canonical arithmetic trivialization — Intro, §0.5, 2.3.2
- canonical coordinates (for a real analytic Kähler metric) — Intro, §0.2
- canonical curve — Intro, §2.1, 2.2
- canonical Galois representation — Intro, §0.9, 1.1, 1.7
- canonical height — Intro, §1.3; I, §3.2
- canonical indigenous bundle (over \mathbf{C}) — Intro, §0.3
- canonical (multi-)uniformizing p -divisible group — VIII, §2.2
- canonical representation (of a Riemann surface) — Intro, §0.3

- canonical system of modular Frobenius liftings — Intro, §1.1, 1.5; VII, §1.3
- canonical weak uniformizing \mathcal{MF}^∇ -object, (log) p -divisible group IX, §2.1
- classical indigenous bundle — Intro, §1.2; I, §4.1
- classical ordinary node — Intro, §1.2; V, §1
- classical ordinary theory — Intro, §0.9
- connectedness of the moduli stack of curves — Intro, §1.1, 1.4
- conservatively spiked (bundle) — V, §2.1
- critical niche (inner, outer) — V, §1
- crys-stable bundle — Intro, §1.3; I, §1.2
- crystalline contractible — Intro, §1.1, 1.4
- crystalline induction — Intro, §0.9, 1.1, 1.7; IX, §2.3; X, §0

- deperfection — VII, §2.1
- diameter — I, §1.2
- differential p -divisible group — VIII, §2.1
- dormant — Intro, §1.1, 1.2; II, §1.1, 2.1
- dormant for Π — III, §1.1

- epiperfect — Intro, §1.4; VI, §1.1
- F -bundle — Appendix, §3.2
- formal uniformizing \mathcal{MF}^∇ -object — IX, §1.5
- Frobenius action on the Schwarz torsor — Intro, §0.6
- Fuchsian uniformization — Intro, §0.1

- Galois mantle — Intro, §1.7; IX, §2.3, 2.5
- generalized ordinary theory — Intro, §1.1, 1.5
- geometrization — Intro, §1.1, 1.6; VIII, §1.2
- geometrizing subbundle — IX, §1.4
- geometrizing sub-object — IX, §2.2
- grafted node — Intro, §1.2; V, §1
- Grothendieck's anabelian philosophy — Intro, §0.10

- Hodge section — Intro, §0.3, 1.3; I §3.2
- Hodge structure — Intro, §1.3
- Hodge subspace (for a spiked Frobenius lifting) — IX, §1.5
- home — Intro, §1.4; III, §1.1
- hyperbolic Riemann surface of finite type — Intro, §0.1

- indigenization — II, §1.4
- indigenized curve — V, §0
- indigenous bundle — Intro, §0.3
- indigenous for Π — III, §1.1
- infinite Weil restriction of scalars — VI, §1.2
- integrality (in the Arakelov sense) — Intro, §0.4
- intrinsic Hodge theory — Intro, §0.10
- Kodaira-Spencer morphism — Intro, §0.3

Kummer theory — VIII, §2.1

left ordering — VI, §2.4

level — I, §3.2

link — III, §1.1

link stack — III, §1.2

log-curve (stable) — I, §1.1

Lubin-Tate geometry — Intro, §1.6; VIII, §2.5

Lubin-Tate stack — IV, §2.3

mantle — IX, §1.3

\mathcal{MF}^∇ -object — Intro, §1.3

mildly spiked (of strength d) — II, §3.1

molecule — Intro, §1.1, 1.2; V, §0

monodromy (group, question) — Intro, §2.1

multi-canonical point — VIII, §2.3

multi-module — V, §1

multi-uniformization — VIII, §2.3

Mumford's uniformization — Intro, §0.1

naive Frobenius pull-back — VI, §1.4

n -connection — II, §2.1

niche — V, §1

nilcurve — Intro, §0.9; V, §0

nilindigenized curve — V, §0

nilpotent — II, §1.1, 2.1

ω -closed — VII, §2.3

ordinary locus — VII, §1.2

ordinary pattern — VII, §1.4

ordinary (stable) bundle — Intro, §1.8; Appendix, §3.2

outer action — X, §1.4

padding degree — Intro, §1.2; V, §1

p -adic Teichmüller theory — Intro, §0

pants — Intro, §1.2; V, §0

perfection — VI, §1.1

philial node — Intro, §1.2; V, §1

Π -adjacent — III, §1.1

Π -ind-adjacent — III, §1.1

Π -indigenous bundle — Intro, §1.4; VI, §1.4

Π -ordinary — VII, §1.1

plot — Intro, §1.2; V, §1

p^n -curvature — II, §2.1

pre-home — Intro, §1.4; III, §1.1

pre- n -connection — II, §2.1

preplot — V, §1

product of signs — V, §1

- pseudo-torally crys-stable bundle — II, §1.6
- pure tone — Intro, §1.4; IV, §2.3
- QF-canonical curve — Intro, §2.2.4
- quasi-affine stack — III, §2.3
- quasi-Fuchsian group — Intro, §0.4, 2.2.1
- radimmersion — I, §2.4
- radius — Intro, §1.2; I, §1.2
- renormalized Frobenius pull-back — Intro, §0.9, 1.4; VI, §1.4
- renormalized Frobenius pull-back of the mantle — IX, §1.4
- restrictable type — I, §3.3
- right ordering — VI, §2.4
- Rigidity Theorem — Intro, §2.2.3
- scenario — Intro, §1.2; V, §1
- Schottky uniformization — Intro, §0.1
- Schwarz torsor — Intro, §0.4; I, §4.3
- set-theoretic canonical Galois representation — VIII, §1.1;
IX, §2.1
- shifted VF-stack — III, §1.3
- shifting permutation — VIII, §1.1
- Shimura curve — Intro, §0.9, 2.2.4; VIII, §2.5; X, §1.3
- σ -canonical connection — VIII, §2.1
- σ -canonical differential/tangential local system — VIII, §2.1
- σ -canonical lifting, point — VIII, §1.1
- sign at a niche — V, §1
- special — Appendix, §1.1
- spiked geometry — Intro, §1.6; IX, §2.4
- spiked locus — II, §3.1
- spiked nilcurve — Intro, §1.1, 1.2
- spiked VF-pattern — IV, §3.1
- spk-indigenous — VII, §3.3
- stable — Appendix, §1.1
- stack of quasi-analytic self-isogenies — Intro, §1.4, 1.5; VI, §1.4
- standard monodromy endomorphism — I, §4.2
- stratification of $\mathcal{N}_{g,r}$ — Intro, §1.2
- strictly weak pair of Frobenius liftings — IX, §2.1
- strong differential local system — IX, §1.2
- strong perfection — IX, §2.1
- strong portion of the mantle — IX, §1.3
- strong set-theoretic canonical Galois representation — IX, §2.1
- strong uniformization — IX, §1.2
- strong uniformizing p -divisible group — IX, §1.2
- supersingular pattern — VII, §1.4
- system of Frobenius liftings — VIII, §1.1
- taut — V, §1




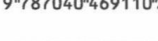











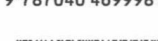






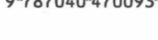



- Teichmüller group (profinite) — X, §1.4
- “third real dimension” (Frobenius as) — Intro, §2.2.6
- torally admissible — II, §1.1
- torally crys-stable bundle — I, §1.2
- torally indigenous bundle — Intro, §1.2; I, §4.1
- torally ordinary — V, §2.1
- total-ordinary rank — V, §2.1
- Torelli map — Intro, §0.9
- total Frobenius lifting — VIII, §1.1
- type of a Riemann surface — Intro, §0.1




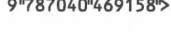














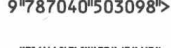





- universal PD-thickening — VI, §1.1
- versal at infinity — II, §2.3
- Verschiebung morphism — II, §1.3
- very ordinary locus — VII, §3.2
- very ordinary spiked Frobenius lifting (of colevel c) — IX, §1.1
- VF-pattern of period ϖ — Intro, §1.1, 1.4; III, §1.1
- VF-stack — III, §1.3
- virtual p -curvature — Intro, §1.2; V, §2.2
- weak differential local system — IX, §1.2
- Weil Conjectures — Appendix, §3.4
- Weil-Petersson metric — Intro, §0.2

郑重声明

高等教育出版社依法对本书享有专有出版权。任何未经许可的复制、销售行为均违反《中华人民共和国著作权法》，其为人将承担相应的民事责任和行政责任；构成犯罪的，将被依法追究刑事责任。为了维护市场秩序，保护读者的合法权益，避免读者误用盗版书造成不良后果，我社将配合行政执法部门和司法机关对违法犯罪的单位和个人进行严厉打击。社会各界人士如发现上述侵权行为，希望及时举报，本社将奖励举报有功人员。

反盗版举报电话	(010) 58581999 58582371 58582488
反盗版举报传真	(010) 82086060
反盗版举报邮箱	dd@hep.com.cn
通信地址	北京市西城区德外大街 4 号 高等教育出版社法律事务与版权管理部
邮政编码	100120

- 1 **Lars V. Ahlfors**, Lectures on Quasiconformal Mappings, Second Edition 
- 2 **Dmitri Burago, Yuri Burago, Sergei Ivanov**, A Course in Metric Geometry 
- 3 **Tobias Holck Colding, William P. Minicozzi II**,
A Course in Minimal Surfaces 
- 4 **Javier Duoandikoetxea**, Fourier Analysis 
- 5 **John P. D'Angelo**, An Introduction to Complex Analysis and Geometry 
- 6 **Y. Eliashberg, N. Mishachev**, Introduction to the b -Principle 
- 7 **Lawrence C. Evans**, Partial Differential Equations, Second Edition 
- 8 **Robert E. Greene, Steven G. Krantz**,
Function Theory of One Complex Variable, Third Edition 
- 9 **Thomas A. Ivey, J. M. Landsberg**,
Cartan for Beginners: Differential Geometry via Moving Frames and
Exterior Differential Systems 
- 10 **Jens Carsten Jantzen**, Representations of Algebraic Groups, Second Edition 
- 11 **A. A. Kirillov**, Lectures on the Orbit Method 
- 12 **Jean-Marie De Koninck, Armel Mercier**,
1001 Problems in Classical Number Theory 
- 13 **Peter D. Lax, Lawrence Zalcman**, Complex Proofs of Real Theorems 
- 14 **David A. Levin, Yuval Peres, Elizabeth L. Wilmer**,
Markov Chains and Mixing Times 
- 15 **Dusa McDuff, Dietmar Salamon**,
 J -holomorphic Curves and Symplectic Topology 
- 16 **John von Neumann**, Invariant Measures 
- 17 **R. Clark Robinson**, An Introduction to Dynamical Systems:
Continuous and Discrete, Second Edition 
- 18 **Terence Tao**, An Epsilon of Room, I: Real Analysis:
pages from year three of a mathematical blog 
- 19 **Terence Tao**, An Epsilon of Room, II:
pages from year three of a mathematical blog 
- 20 **Terence Tao**, An Introduction to Measure Theory 
- 21 **Terence Tao**, Higher Order Fourier Analysis 
- 22 **Terence Tao**, Poincaré's Legacies,
Part I: pages from year two of a mathematical blog 
- 23 **Terence Tao**, Poincaré's Legacies,
Part II: pages from year two of a mathematical blog 
- 24 **Cédric Villani**, Topics in Optimal Transportation 
- 25 **R. J. Williams**, Introduction to the Mathematics of Finance 
- 26 **T. Y. Lam**, Introduction to Quadratic Forms over Fields 

- 27 **Jens Carsten Jantzen**, Lectures on Quantum Groups  9 787040 469141 >
- 28 **Henryk Iwaniec**, Topics in Classical Automorphic Forms  9 787040 469134 >
- 29 **Sigurdur Helgason**, Differential Geometry, Lie Groups, and Symmetric Spaces  9 787040 469165 >
- 30 **John B. Conway**, A Course in Operator Theory  9 787040 469158 >
- 31 **James E. Humphreys**, Representations of Semisimple Lie Algebras in the BGG Category O  9 787040 468984 >
- 32 **Nathanial P. Brown, Narutaka Ozawa**, C^* -Algebras and Finite-Dimensional Approximations  9 787040 469325 >
- 33 **Hiraku Nakajima**, Lectures on Hilbert Schemes of Points on Surfaces  9 787040 501216 >
- 34 **S. P. Novikov, I. A. Taimanov**, Translated by Dmitry Chibisov, Modern Geometric Structures and Fields  9 787040 469189 >
- 35 **Luis Caffarelli, Sandro Salsa**, A Geometric Approach to Free Boundary Problems  9 787040 469202 >
- 36 **Paul H. Rabinowitz**, Minimax Methods in Critical Point Theory with Applications to Differential Equations  9 787040 502299 >
- 37 **Fan R. K. Chung**, Spectral Graph Theory  9 787040 502305 >
- 38 **Susan Montgomery**, Hopf Algebras and Their Actions on Rings  9 787040 502312 >
- 39 **C. T. C. Wall**, Edited by A. A. Ranicki, Surgery on Compact Manifolds, Second Edition  9 787040 502329 >
- 40 **Frank Sottile**, Real Solutions to Equations from Geometry  9 787040 501513 >
- 41 **Bernd Sturmfels**, Gröbner Bases and Convex Polytopes  9 787040 503081 >
- 42 **Terence Tao**, Nonlinear Dispersive Equations: Local and Global Analysis  9 787040 503050 >
- 43 **David A. Cox, John B. Little, Henry K. Schenck**, Toric Varieties  9 787040 503098 >
- 44 **Luca Capogna, Carlos E. Kenig, Loredana Lanzani**, Harmonic Measure: Geometric and Analytic Points of View  9 787040 503074 >
- 45 **Luis A. Caffarelli, Xavier Cabré**, Fully Nonlinear Elliptic Equations  9 787040 503067 >
- 46 **Teresa Crespo, Zbigniew Hajto**, Algebraic Groups and Differential Galois Theory  9 787040 510133 >
- 47 **Barbara Fantechi, Lothar Göttsche, Luc Illusie, Steven L. Kleiman, Nitin Nitsure, Angelo Vistoli**, Fundamental Algebraic Geometry: Grothendieck's FGA Explained  9 787040 510126 >
- 48 **Shinichi Mochizuki**, Foundations of p -adic Teichmüller Theory  9 787040 510089 >
- 49 **Manfred Leopold Einsiedler, David Alexandre Ellwood, Alex Eskin, Dmitry Kleinbock, Elon Lindenstrauss, Gregory Margulis, Stefano Marmi, Jean-Christophe Yoccoz**, Homogeneous Flows, Moduli Spaces and Arithmetic  9 787040 510096 >
- 50 **David A. Ellwood, Emma Previato**, Grassmannians, Moduli Spaces and Vector Bundles  9 787040 510393 >
- 51 **Jeffery McNeal, Mircea Mustață**, Analytic and Algebraic Geometry: Common Problems, Different Methods  9 787040 510553 >
- 52 **V. Kumar Murty**, Algebraic Curves and Cryptography  9 787040 510386 >
- 53 **James Arthur, James W. Cogdell, Steve Gelbart, David Goldberg, Dinakar Ramakrishnan, Jiu-Kang Yu**, On Certain L -Functions  9 787040 510409 >

本书为 p 进双曲线及其模空间的单值化理论奠定了基础。一方面, 这个理论将复双曲线及其模空间的 Fuchs 和 Bers 单值化推广到了非阿基米德情形, 该理论在本书中简称为 p 进 Teichmüller 理论。另一方面, 该理论可以看作是常阿贝尔簇及其模空间的 Serre–Tate 理论的相当精确的双曲模拟。

p 进双曲线及其模空间的单值化理论始于作者先前的一些工作。从某种意义上说, 本书是对先前工作的概括和延续。本书旨在填补所提出方法与在本科复分析课程中研究的双曲黎曼曲面的经典单值化之间的缺口。

- 介绍从 p 进伽罗瓦表示的角度对曲线模空间的一种系统化处理。
- 给出 Serre–Tate 理论的双曲线模拟。
- 建立 Fuchs 和 Bers 单值化理论的 p 进模拟。
- 提供 p 进 Hodge 理论的一个“非阿贝尔例子”的系统化处理。

本版只限于中华人民共和国境内发行。本版经由美国数学会授权仅在中华人民共和国境内销售, 不得出口。

美国数学会经典影印系列



ISBN 978-7-04-051008-9



9 787040 510089 >

定价 199.00 元